Volume I: NeXus User Manual

Copyright © 1996-2012 NeXus International Advisory Committee, http://www.nexusformat.org

The NeXus manual is licensed under the terms of the GNU Free Documentation License version 1.3. See the FDL license file included with the source of this manual or refer to http://www.gnu.org/licenses/fdl-1.3.txt for more details.

The examples in the NeXus manual are licensed under the terms of the GNU Lesser General Public License version 3. See the LGPL license file included with the source of this manual or refer to http://www.gnu.org/licenses/lgpl-3.0.txt for more details.

COLLABORATORS

	TITLE : Volume I: NeXus User Manual	

ACTION	*NAME*	*DATE*	*SIGNATURE*
WRITTEN BY	Ray Osborn, Mark Koennecke, Przemek Klosowski, Frederick Akeroyd, Peter F. Peterson, Pete R. Jemian, Stuart I. Campbell, and Tobias Richter	October 3, 2012	

REVISION HISTORY

NUMBER	DATE	DESCRIPTION	NAME
	2009-11	Started conversion from the old NeXus mediawiki documentation.	PFP
initial draft	2010 spring	Most of the content from the old NeXus mediawiki documentation is included. Some new wiki content has been introduced but should be easy to identify for inclusion in the manual.	PRJ
draft	2010-11	Nearly complete but still much finishing work remains. The description of dimensions and the description of the coordinate system needs major revision and improvement. More examples are needed. The manual is now divided into two volumes. Volume I is the User Manual, Volume II is the Reference Documentation. Much of the NXDL chapter in Volume II is autogenerated from the `nxdl.xsd` Schema and the NXDL source files. Initial release of NXDL, manual, and next release of NAPI (compatibility release), manual in editing stage, packaging tools need to be built.	PRJ
1.0b	2011-11	Preparing manual for initial release. Also preparing to convert manual source from DocBook to Sphinx for next release of manual.	PRJ

REVISION HISTORY			
NUMBER	DATE	DESCRIPTION	NAME
3.1	2012-02	Now using cmake to control multiplatform build and packaging.	PRJ
3.1	2012-05	Ready for release.	PRJ

Contents

1 NeXus Introduction **1**
1.1 What is NeXus? . . . 1
1.1.1 A Set of Design Principles . . . 2
1.1.1.1 Example of a NeXus File . . . 2
1.1.1.2 Important Classes . . . 2
1.1.1.3 Simple Example . . . 3
1.1.2 A Set of Data Storage Objects . . . 5
1.1.3 A Set of Subroutines . . . 5
1.1.4 Scientific Community . . . 6
1.2 Motivations for the NeXus standard in the Scientific Community . . . 6
1.2.1 Simple plotting . . . 7
1.2.2 Unified format for reduction and analysis . . . 7
1.2.2.1 NeXus as a Common Data Exchange Format . . . 7
1.2.3 Defined dictionary of terms . . . 8
1.3 NAPI: The NeXus Application Programming Interface . . . 8
1.3.1 How do I write a NeXus file? . . . 9
1.3.2 How do I read a NeXus file? . . . 10
1.3.3 How do I browse a NeXus file? . . . 10

2 NeXus Design **12**
2.1 NeXus Objects and Terms . . . 12
2.1.1 Data Groups . . . 12
2.1.2 Data Fields . . . 13
2.1.3 Data Attributes . . . 13
2.1.4 Links . . . 14
2.1.5 NeXus Base Classes . . . 14
2.1.5.1 `NXdata` Facilitates Automatic Plotting . . . 15
2.1.6 NeXus Application Definitions . . . 16
2.2 NeXus Coordinate Systems . . . 16
2.2.1 McStas and `NXgeometry` System . . . 16

2.2.2 Simple (Spherical Polar) Coordinate System . . . 17
2.2.3 Coordinate Transformations . . . 18
2.3 Rules for Structuring Information in NeXus Files . . . 19
2.3.1 Content of a Raw Data `NXentry` Group . . . 19
2.3.2 Content of a processed data `NXentry` group . . . 19
2.3.3 `NXsubentry` or Multi-Method Data . . . 20
2.3.4 Rules for Special Cases . . . 21
2.3.4.1 Scans . . . 21
2.3.4.1.1 Simple scan . . . 21
2.3.4.1.2 Simple scan with area detector . . . 22
2.3.4.1.3 Complex *hkl* scan . . . 22
2.3.4.1.4 Multi-parameter scan: XAS . . . 24
2.3.4.2 Rastering . . . 24
2.3.4.3 NXcollection . . . 25
2.4 Rules for Storing Data Items in NeXus Files . . . 25
2.4.1 Naming Conventions . . . 25
2.4.2 NeXus Array Storage Order . . . 26
2.4.2.1 Non C Storage Order . . . 26
2.4.3 NeXus Data Types . . . 28
2.4.3.1 NeXus dates and times . . . 29
2.4.4 NeXus Data Units . . . 30
2.4.5 Linking Multi Dimensional Data with Axis Data . . . 30
2.4.5.1 Linking by name using the `axes` attribute . . . 30
2.4.5.2 Linking by dimension number using the `axis` attribute . . . 30
2.4.5.3 Discussion of the two linking methods . . . 31
2.4.6 Storing Detectors . . . 31
2.4.7 Monitors are Special . . . 32
2.4.8 Find the plottable data . . . 32
2.5 Physical File format . . . 32
2.5.1 Choice of HDF as Underlying File Format . . . 32
2.5.2 Mapping NeXus into HDF . . . 33
2.5.3 Mapping NeXus into XML . . . 35
2.5.4 Special Attributes . . . 36

3 Constructing NeXus Files and Application Definitions **37**
3.1 The WOnderful New Instrument (WONI) . . . 37
3.2 Constructing a NeXus file for WONI . . . 39
3.2.1 Decide which parameters need to be stored . . . 39
3.2.2 Mapping parameters to NeXus . . . 39
3.2.3 Decide on `NXdata` . . . 40
3.2.4 Fill in auxiliary Information . . . 40
3.3 Creating a NXDL Specification . . . 40
3.3.1 Application Definition Steps . . . 40
3.3.2 Step 1: *Think!* hard about data . . . 41
3.3.3 Step 2: *Map* Data into the NeXus Hierarchy . . . 41
3.3.4 Step 3: *Describe* this map in a NXDL file . . . 43
3.3.5 Step 4: *Standardize* with the NIAC . . . 44
3.3.6 Full listing of the WONI Application Definition . . . 44
3.3.7 Using an Application Definition . . . 47
3.4 Processed Data . . . 47

4 Brief history of the NeXus format **48**

5 NeXus Community **50**
5.1 NIAC: The NeXus International Advisory Committee . . . 50
5.2 NeXus Mailing Lists . . . 51
5.3 NeXus Subversion Repositories . . . 52
5.3.1 Login . . . 52
5.3.2 Committing Changes . . . 53
5.3.3 URLs described in this section . . . 53
5.4 NeXus Issue Reporting . . . 54
5.4.1 NeXus Code (Library and Applications) . . . 55
5.4.2 NeXus Definitions (NXDL base classes and application definitions) . . . 55

A Installation **56**
A.1 Precompiled Binary Installation . . . 56
A.1.1 Prerequisites . . . 56
A.1.1.1 HDF5/HDF4 . . . 56
A.1.1.2 XML . . . 56
A.1.2 Linux RPM Distribution Kits . . . 56
A.1.3 Microsoft Windows Installation Kit . . . 57
A.1.4 Mac OS X Installation Kit . . . 57
A.2 Source Installation . . . 57
A.2.1 NeXus Source Code Distribution . . . 57
A.2.2 Cygwin Kits . . . 58

B Verification and validation of files **59**
B.1 Overview . . . 59
B.2 Definitions of these terms . . . 60
B.3 NeXus data files may use multiple base classes or application definitions . . . 61
B.4 Validation techniques . . . 61
B.4.1 Validation of NeXus data files . . . 61
B.4.2 Validation of NeXus Definition Language (NXDL) specification files . . . 61
B.4.3 Validation of the NXDL rules . . . 62
B.4.4 Validation of XSLT files . . . 62
B.4.5 Transformation of NXDL files to Schematron . . . 62

C NeXus Utilities **63**
C.1 Utilities supplied with NeXus . . . 63
C.2 Data Analysis . . . 64
C.3 HDF Tools . . . 64

D Frequently Asked Questions **66**

E Licenses **68**
E.1 FDL: GNU Free Documentation License . . . 68
E.2 LGPL: GNU Lesser Gnu Public License . . . 75

6 Index **79**

List of Figures

1 example of NeXus data x

1.1 Example of a NeXus file 2
1.2 NXinstrument excerpt 3
1.3 Simple Example 3
1.4 *N* separate file formats 8
1.5 *N* separate file formats joined by a common NeXus converter 8

2.1 Linking in a NeXus file 14
2.2 The McStas Coordinate System 17
2.3 NeXus Simple (Spherical Polar) Coordinate System 18
2.4 NeXus Raw Data Hierarchy 19
2.5 NeXus Processed Data Hierarchy 20
2.6 NeXus Multi Method Hierarchy 21
2.7 NeXus Simple Scan Example 22
2.8 NeXus Simple Scan Example with Area Detector 22
2.9 NeXus Complex *hkl* Scan 23
2.10 NeXus Multi-parameter scan: XAS 24
2.11 NXcollection Example 25

3.1 The (fictional) WONI example powder diffractometer 38
3.2 Example Powder Diffraction Plot from (fictional) WONI at HYNES 38
3.3 Basic structure of a NeXus file 39
3.4 Basic structure of a NeXus file with a monochromator added 39

B.1 Flowchart of the NeXus validation process. 60

Preface

With this edition of the manual, NeXus introduces a complete version of the documentation of the NeXus standard. The content from the wiki has been converted, augmented (in some parts significantly), clarified, and indexed. The NeXus Definition Language (NXDL) is introduced now to define base classes and application definitions. NXDL replaces the previous method (meta-DTD) to define NeXus classes. NeXus base classes and instrument definitions are now assigned to one of three classifications: (1) *base classes* (that represent the components used to build a NeXus data file), (2) *application definitions* (used to define a minimum set of data for a specific purpose such as scientific data processing or an instrument definition), and (3) *contributed definitions* (definitions and specifications that are in an incubation status before ratification by the NIAC). Additional examples have been added to respond to inquiry from the users of the NeXus standard about implementation and usage. Hopefully, the improved documentation with more examples and the new NXDL will reduce the learning barriers incurred by those new to NeXus.

Representation of data examples

Most of the examples of data files have been written in a format intended to show the structure of the file rather than the data content. In some cases, where it is useful, some of the data is shown. Consider this prototype example:

```
entry:NXentry
  instrument:NXinstrument
    detector:NXdetector
      data:[]
        @axes = "bins"
        @long_name = "strip detector 1-D array"
        @signal = 1
      bins:[0, 1, 2, ... 1023]
        @long_name = "bin index numbers"
  sample:NXsample
    name = "zeolite"
  data:NXdata
    data --> /entry/instrument/detector/data
    bins --> /entry/instrument/detector/bins
```

Figure 1: example of NeXus data

Some words on the notation:

- Hierarchy is represented by indentation. Objects on the same indentation level are in the same group
- The combination `name:NXclass` denotes a NeXus group with name `name` and class `NXclass`.
- A simple name (no following class) denotes a data field. An equal sign is used to show the value, where this is important to the example.
- Sometimes, a data type is specified and possibly a set of dimensions. For example, `energy:NX_NUMBER[NE]` says *energy* is a 1-D array of numbers (either integer or floating point) of length `NE`.
- Attributes are noted as @name="value" pairs. The `@` symbol only indicates this is an attribute and is not part of the attribute name.
- Links are shown with a text arrow `-->` indicating the source of the link (using HDF5 notation listing the sequence of *names*).

Line 1 shows that there is one group at the root level of the file named `entry`. This group is of type `NXentry` which means it conforms to the specification of the `NXentry` NeXus base class. Using the HDF5 nomenclature, we would refer to this as the `/entry` group.

Lines 2, 10, and 12: The `/entry` group contains three subgroups: `instrument`, `sample`, and `data`. These groups are of type `NXinstrument`, `NXsample`, and `NXdata`, respectively.

Line 4: The data of this example is stored in the `/entry/instrument/detector` group in the dataset called `data` (HDF5 path is `/entry/instrument/detector/data`). The indication of `data:[]` says that `data` is an array of unspecified dimension(s).

Lines 5-7: There are three attributes of `/entry/instrument/detector/data`: `axes`, `long_name`, and `signal`.

Line 8 (reading `bins:[0, 1, 2, ... 1023]`) shows that `bins` is a 1-D array of length presumably 1024. A small, representative selection of values are shown.

Line 9: an attribute that shows a descriptive name of `/entry/instrument/detector/bins`. This attribute might be used by a NeXus client while plotting the data.

Line 11 (reading `name = "zeolite"`) shows how a string value is represented.

Lines 13-14: The `/entry/data`) group has two datasets that are actually linked as shown. (As you will see later, the `NXdata` group is required and enables NeXus clients to easily determine what to offer for display on a default plot.)

Class path specification

In some places in this documentation, a path may be shown using the class types rather than names. For example: `/NXentry/NXinstrument/NXcrystal/wavelength` identifies a dataset called `wavelength` that is inside a group of type `NXcrystal` ... This nomemclature is used when the exact name of each group is either unimportant or not specified. Often, this will be used in a NXDL specification to indicate the connections of a link.

Chapter 1

NeXus Introduction

In recent years, a community of scientists and computer programmers working in neutron and synchrotron facilities around the world came to the conclusion that a common data format would fulfill a valuable function in the scattering community. As instrumentation becomes more complex and data visualization become more challenging, individual scientists, or even institutions, have found it difficult to keep up with new developments. A common data format makes it easier, both to exchange experimental results and to exchange ideas about how to analyze them. It promotes greater cooperation in software development and stimulates the design of more sophisticated visualization tools. Additional background information is given in A Brief History of the NeXus Format.

This section is designed to give a brief introduction to NeXus, the data format and tools that have been developed in response to these needs. It explains what a modern data format such as NeXus is and how to write simple programs to read and write NeXus files.

The programmers who produce intermediate files for storing analyzed data should agree on simple interchange rules.

1.1 What is NeXus?

The NeXus data format has four components:

A set of design principles to help people understand what is in the data files.

A set of data storage objects (Base Classes and Application Definitions) to allow the development of portable analysis software.

A set of subroutines (NeXus Utilities) to make it easy to read and write NeXus data files.

A Scientific Community to provide the scientific data, advice, and continued involvement with the NeXus standard. NeXus provides a forum for the scientific community to exchange ideas in data storage.

In addition, NeXus relies on a set of low-level file formats to actually store NeXus files on physical media. Each of these components are described in more detail in the Physical File Format section.

The NeXus Application-Programmer Interface (NAPI), which provides the set of subroutines for reading and writing NeXus data files, is described briefly in Section 1.3. (Further details are provided in the NAPI chapter of Volume II of this documentation.) The principles guiding the design and implementation of the NeXus standard are described in the NeXus Design chapter. Base classes, which comprise the data storage objects used in NeXus data files, are detailed in the Base Classes chapter of Volume II of this documentation. Additionally, a brief list describing the set of NeXus Utilities available to browse, validate, translate, and visualise NeXus data files is provided in the NeXus Utilities chapter.

1.1.1 A Set of Design Principles

NeXus data files contain four types of entity: data groups, data fields, attributes, and links.

Data Groups Data groups are like folders that can contain a number of fields and/or other groups.

Data Fields Data fields can be scalar values or multidimensional arrays of a variety of sizes (1-byte, 2-byte, 4-byte, 8-byte) and types (characters, integers, floats). In HDF, fields are represented as HDF *Scientific Data Sets* (also known as SDS).

Data Attributes Extra information required to describe a particular group or field, such as the data units, can be stored as a data attribute.

Links Links are used to reference the plottable data from `NXdata` when the data is provided in other groups such as `NXmonitor` or `NXdetector`.

In fact, a NeXus file can be viewed as a computer file system. Just as files are stored in folders (or subdirectories) to make them easy to locate, so NeXus fields are stored in groups. The group hierarchy is designed to make it easy to navigate a NeXus file.

1.1.1.1 Example of a NeXus File

The following diagram shows an example of a NeXus file represented as a tree structure.

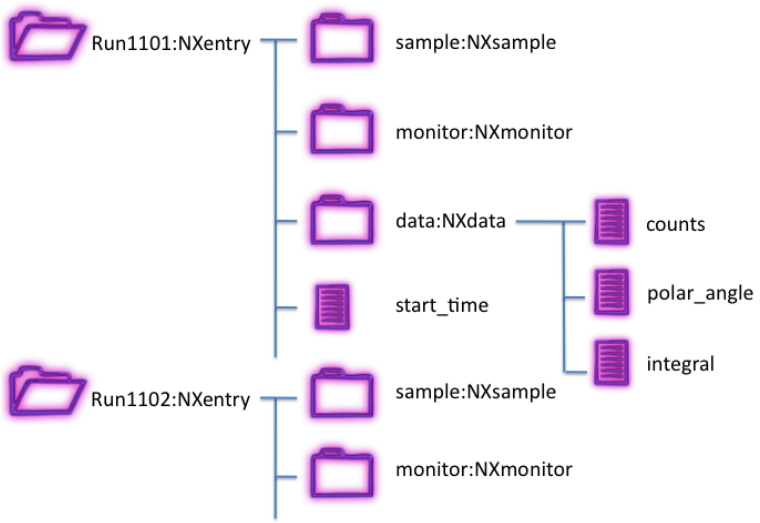

Figure 1.1: Example of a NeXus file

Note that each field is identified by a name, such as `counts`, but each group is identified both by a name and, after a colon as a delimiter, the class type, e.g., `monitor:NXmonitor`). The class types, which all begin with `NX`, define the sort of fields that the group should contain, in this case, counts from a beamline monitor. The hierarchical design, with data items nested in groups, makes it easy to identify information if you are browsing through a file.

1.1.1.2 Important Classes

Here are some of the important classes found in nearly all NeXus files. A complete list can be found in the NeXus Design chapter.

Note
Note that `NXentry` and `NXdata` are the only two classes necessary to store the minimum amount of information in a valid NeXus data file.

NXentry **Required:** The top level of any NeXus file contains one or more groups with the class `NXentry`. These contain all the data that is required to describe an experimental run or scan. Each `NXentry` typically contains a number of groups describing sample information (class `NXsample`), instrument details (class `NXinstrument`), and monitor counts (class `NXmonitor`).

NXdata **Required:** Each `NXentry` group contains one or more groups with class `NXdata`. These groups contain the experimental results in a self-contained way, i.e., it should be possible to generate a sensible plot of the data from the information contained in each `NXdata` group. That means it should contain the axis labels and titles as well as the data.

NXsample A `NXentry` group will often contain a group with class `NXsample`. This group contains information pertaining to the sample, such as its chemical composition, mass, and environment variables (temperature, pressure, magnetic field, etc.).

NXinstrument There might also be a group with class `NXinstrument`. This is designed to encapsulate all the instrumental information that might be relevant to a measurement, such as flight paths, collimations, chopper frequencies, etc.

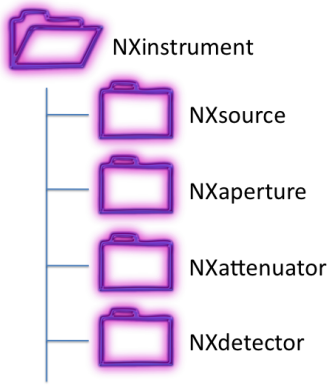

Figure 1.2: NXinstrument excerpt

Since an instrument can comprise several beamline components each defined by several parameters, they are each specified by a separate group. This hides the complexity from generic file browsers, but makes the information available in an intuitively obvious way if it is required.

1.1.1.3 Simple Example

NeXus data files do not need to be complicated. In fact, the following diagram shows an extremely simple NeXus file (in fact, the simple example shows the minimum information necessary for a NeXus data file) that could be used to transfer data between programs. (Later in this section, we show how to write and read this simple example.)

Figure 1.3: Simple Example

This illustrates the fact that the structure of NeXus files is extremely flexible. It can accommodate very complex instrumental information, if required, but it can also be used to store very simple data sets. In the next example, a NeXus data file is shown as XML:

Example 1.1 *verysimple.xml*: A very simple NeXus Data file (in XML)

```
<?xml version="1.0" encoding="UTF-8"?>
  <NXroot NeXus_version="4.3.0" XML_version="mxml"
      file_name="verysimple.xml"
      xmlns="http://definition.nexusformat.org/schema/3.1"
      xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
      xsi:schemaLocation="http://definition.nexusformat.org/schema/3.1
                          http://definition.nexusformat.org/schema/3.1/BASE.xsd"
      file_time="2010-11-12T12:40:17-06:00">
    <NXentry name="entry">
      <NXdata name="data">
        <counts NAPItype="NX_INT64[15]" long_name="photodiode counts" signal="NX_INT32:1"
            axes="two_theta">
                              1193                    4474
                             53220                  274310
                            515430                  827880
                           1227100                 1434640
                           1330280                 1037070
                            598720                  316460
                             56677                    1000
                              1000
        </counts>
        <two_theta NAPItype="NX_FLOAT64[15]" units="degrees" long_name="two_theta (degrees)
            ">
                  18.90940         18.90960         18.90980         18.91000
                  18.91020         18.91040         18.91060         18.91080
                  18.91100         18.91120         18.91140         18.91160
                  18.91180         18.91200         18.91220
        </two_theta>
      </NXdata>
    </NXentry>
  </NXroot>
```

NeXus files are easy to create. This example NeXus file was created using a short Python program and NeXpy:

Example 1.2 *verysimple.py*: Using NeXpy to write a very simple NeXus Data file (in HDF5)

```
#
# This example uses NeXpy to build the verysimple.nx5 data file.

from nexpy.api import nexus

angle = [18.9094, 18.9096, 18.9098, 18.91,  18.9102,
         18.9104, 18.9106, 18.9108, 18.911, 18.9112,
   18.9114, 18.9116, 18.9118, 18.912, 18.9122]
diode = [1193, 4474, 53220, 274310, 515430, 827880,
         1227100, 1434640, 1330280, 1037070, 598720,
   316460, 56677, 1000, 1000]

two_theta = nexus.SDS(angle, name="two_theta",
                units="degrees",
         long_name="two_theta (degrees)")
counts = nexus.SDS(diode, name="counts", long_name="photodiode counts")
data = nexus.NXdata(counts,[two_theta])
data.nxsave("verysimple.nx5")

# The verysimple.xml file was built with this command:
# nxconvert -x verysimple.nx5 verysimple.xml
# and then hand-edited (line breaks) for display.
```

1.1.2 A Set of Data Storage Objects

If the design principles are followed, it will be easy for anyone browsing a NeXus file to understand what it contains, without any prior information. However, if you are writing specialized visualization or analysis software, you will need to know precisely what specific information is contained in advance. For that reason, NeXus provides a way of defining the format for particular instrument types, such as time-of-flight small angle neutron scattering. This requires some agreement by the relevant communities, but enables the development of much more portable software.

The set of data storage objects is divided into three parts: base classes, application definitions, and contributed definitions. The base classes represent a set of components that define the dictionary of all possible terms to be used with that component. The application definitions specify the minimum required information to satisfy a particular scientific or data analysis software interest. The contributed definitions have been submitted by the scientific community for incubation before they are adopted by the NIAC or for availability to the community.

These instrument definitions are formalized as XML files, using NXDL, (as described in the NXDL chapter in Volume II of this documentation) to specify the names of data fields, and other NeXus data objects. The following is an example of such a file for the simple NeXus file shown above.

Example 1.3 *verysimple.nxdl.xml*: A very simple NeXus Definition Language (NXDL) file

```
<?xml version="1.0" ?>
<definition
  xmlns="http://definition.nexusformat.org/nxdl/3.1"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://definition.nexusformat.org/nxdl/3.1 ../nxdl.xsd"
  category="base"
  name="verysimple"
  version="1.0"
  svnid="$Id: verysimple.nxdl.xml 730 2010-11-12 18:40:01Z Pete Jemian $"
  type="group" extends="NXobject">

  <doc>
    A very simple NeXus NXDL file
  </doc>
  <group type="NXentry">
    <group type="NXdata">
      <field name="counts" type="NX_INT"  units="NX_UNITLESS">
        <doc>counts recorded by detector</doc>
      </field>
      <field name="two_theta" type="NX_FLOAT" units="NX_ANGLE">
        <doc>rotation angle of detector arm</doc>
      </field>
    </group>
  </group>
</definition>
```

For complete examples of reading and writing NeXus data files, refer to the Examples of reading or writing NeXus data files chapter in Volume II. This chapter has several examples of writing and reading NeXus data files. If you want to define the format of a particular type of NeXus file for your own use, e.g. as the standard output from a program, you are encouraged to *publish* the format using this XML format. An example of how to do this is shown in the section titled Creating a NXDL Specification.

1.1.3 A Set of Subroutines

NeXus data files are high-level so the user only needs to know how the data are referenced in the file but does not need to be concerned where the data are stored in the file. Thus, the data are most easily accessed using a subroutine library tuned to the specifics of the data format.

In the past, a data format was defined by a document describing the precise location of every item in the data file, either as row and column numbers in an ASCII file, or as record and byte numbers in a binary file. It is the job of the subroutine library to retrieve the data. This subroutine library is commonly called an application-programmer interface or API.

For example, in NeXus, a program to read in the wavelength of an experiment would contain lines similar to the following:

Example 1.4 Simple example of reading data using the NeXus API

```
NXopendata (fileID, "wavelength");
NXgetdata (fileID, lambda);
NXclosedata (fileID);
```

In this example, the program requests the value of the data that has the label `wavelength`, storing the result in the variable lambda. `fileID` is a file identifier that is provided by NeXus when the file is opened.

We shall provide a more complete example when we have discussed the contents of the NeXus files.

1.1.4 Scientific Community

NeXus began as a group of scientists with the goal of defining a common data storage format to exchange experimental results and to exchange ideas about how to analyze them.

The NeXus Scientific Community provides the scientific data, advice, and continued involvement with the NeXus standard. NeXus provides a forum for the scientific community to exchange ideas in data storage through the NeXus wiki.

The NeXus International Advisory Committee supervises the development and maintenance of the NeXus common data format for neutron, x-ray, and muon science. The NIAC: The NeXus International Advisory Committee supervises a technical committee to oversee the NeXus Application Programmer Interface (NAPI: The NeXus Application Programming Interface) and the NeXus class definitions.

1.2 Motivations for the NeXus standard in the Scientific Community

By the early 1990s, several groups of scientists in the fields of neutron and X-ray science had recognized a common and troublesome pattern in the data acquired at various scientific instruments and user facilities. Each of these instruments and facilities had a locally defined format for recording experimental data. With lots of different formats, much of the scientists' time was being wasted in the task of writing import readers for processing and analysis programs. As is common, the exact information to be documented from each instrument in a data file evolves, such as the implementation of new high-throughput detectors. Many of these formats lacked the generality to extend to the new data to be stored, thus another new format was devised. In such environments, the documentation of each generation of data format is often lacking.

Three parallel developments have led to NeXus:

1. *June 1994*: Mark Koennecke (Paul Scherer Institute, Switzerland) made a proposal using netCDF for the European neutron scattering community while working at the ISIS pulsed neutron facility.
2. *August 1994*: Jon Tischler and Mitch Nelson (Oak Ridge National Laboratory, USA) proposed an HDF-based format as a standard for data storage at the Advanced Photon Source (Argonne National Laboratory, USA).
3. *October 1996*: Przemek Klosowski (National Institute of Standards and Technology, USA) produced a first draft of the NeXus proposal drawing on ideas from both sources.

These scientists proposed methods to store data using a self-describing, extensible format that was already in broad use in other scientific disciplines. Their proposals formed the basis for the current design of the NeXus standard which was developed across three workshops organized by Ray Osborn (ANL), *SoftNeSS'94* (Argonne Oct. 1994), *SoftNeSS'95* (NIST Sept. 1995), and *SoftNeSS'96* (Argonne Oct. 1996), attended by representatives of a range of neutron and x-ray facilities. The NeXus API was released in late 1997. Basic motivations for this standard were:

BASIC MOTIVATIONS FOR THE NEXUS STANDARD

1. Simple plotting
2. A Unified format for reduction and analysis
3. A Defined dictionary of terms

1.2.1 Simple plotting

An important motivation for the design of NeXus was to simplify the creation of a default plot view. While the best representation of a set of observations will vary, depending on various conditions, a good suggestion is often known *a priori*. This suggestion is described in the `NXdata` element so that any program that is used to browse NeXus data files can provide a *best representation* without request for user input.

1.2.2 Unified format for reduction and analysis

Another important motivation for NeXus, indeed the *raison d'etre*, was the community need to analyze data from different user facilities. A single data format that is in use at a variety of facilities would provide a major benefit to the scientific community. This unified format should be capable of describing any type of data from the scientific experiments, at any step of the process from data acquisition to data reduction and analysis. This unified format also needs to allow data to be written to storage as efficiently as possible to enable use with high-speed data acquisition.

Self-description, combined with a reliance on a **multi-platform** (and thereby **portable**) data storage format, are valued components of a data storage format where the longevity of the data is expected to be longer than the lifetime of the facility at which it is acquired. As the name implies, self-description within data files is the practice where the structure of the information contained within the file is evident from the file itself. A multi-platform data storage format must faithfully represent the data identically on a variety of computer systems, regardless of the bit order or byte order or word size native to the computer.

The scientific community continues to grow the various types of data to be expressed in data files. This practice is expected to continue as part of the investigative process. To gain broad acceptance in the scientific user community, any data storage format proposed as a standard would need to be **extendable** and continue to provide a means to express the latest notions of scientific data.

The maintenance cost of common data structures meeting the motivations above (self-describing, portable, and extendable) is not insurmountable but is often well-beyond the research funding of individual members of the muon, neutron, and X-ray science communities. Since it is these members that drive the selection of a data storage format, it is necessary for the user cost to be as minimal as possible. In this case, experience has shown that the format must be in the **public-domain** for it to be commonly accepted as a standard. A benefit of the public-domain aspect is that the source code for the API is open and accessible, a point which has received notable comment in the scientific literature.

More recently, NeXus has recognized that part of the scientific community with a desire to write and record scientific data, has small data volumes and a large aversion to the requirement of a complicated API necessary to access data in binary files such as HDF. For such information, the NeXus API has been extended by the addition of the eXtensible Markup Language[1] (XML) as an alternative to HDF. XML is a text-based format that supports compression and structured data and has broad usage in business and e-commerce. While possibly complicated, XML files are human readable, and tools for translation and extraction are plentiful. The API has routines to read and write XML data and to convert between HDF and XML.

1.2.2.1 NeXus as a Common Data Exchange Format

By the late 1980s, it had become common practice for a scientific instrument or facility to define its own data format, often at the convenience of the local computer system. Data from these facilities were not easily interchanged due to various differences in computer systems and the compression schemes of binary data. It was necessary to contact the facility to obtain a description so that one could write an import routine in software. Experience with facilities closing (and subsequent lack of access to information describing the facility data format) revealed a significant limitation with this common practice. Further, there existed a $N * N$ number of conversion routines necessary to convert data between various formats. In Figure 1.4, circles represent different data file formats while arrows represent conversion routines. Note that the red circle only maps to one other format.

[1] XML: http://www.w3.org/XML/. There are many other descriptions of XML, for example: http://en.wikipedia.org/wiki/XML

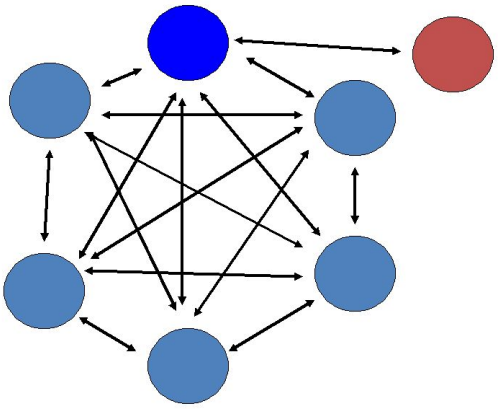

Figure 1.4: *N* separate file formats

One early idea has been for NeXus to become the common data exchange format, and thereby reduce the number of data conversion routines from *N* * *N* down to *2N*, as show in Figure 1.5.

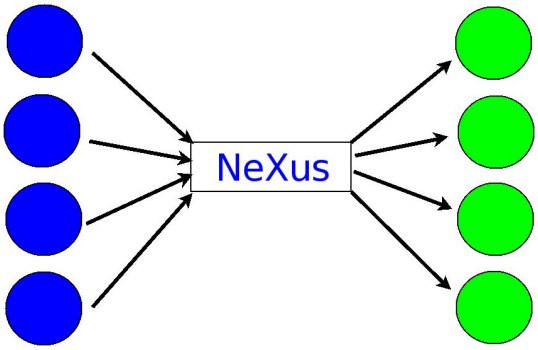

Figure 1.5: *N* separate file formats joined by a common NeXus converter

1.2.3 Defined dictionary of terms

A necessary feature of a standard for the interchange of scientific data is a *defined dictionary* (or *lexicography*) of terms. This dictionary declares the expected spelling and meaning of terms when they are present so that it is not necessary to search for all the variant forms of *energy* when it is used to describe data (e.g., `E`, `e`, `keV`, `eV`, `nrg`, ...).

NeXus recognized that each scientific specialty has developed a unique dictionary and needs to categorize data using those terms. The NeXus Application Definitions provide the means to document the lexicography for use in data files of that scientific specialty.

1.3 NAPI: The NeXus Application Programming Interface

The NeXus API consists of routines to read and write NeXus data files. It was written to provide a simple to use and consistent common interface for all supported backends (XML, HDF4 and HDF5) to scientific programmers and other users of the NeXus Data Standard.

This section will provide a brief overview of the available functionality. Further documentation of the NeXus Application Programming Interface (NAPI) for bindings to specific programming language can be found in the NAPI chapter in Volume II of this documentation and obtained from the NeXus development site.[2]

[2] http://download.nexusformat.org

For an even more detailed description of the internal workings of NAPI see NeXusIntern.pdf[3] in the NeXus code repository. That document is written for programmers who want to work on the NAPI itself. If you are new to NeXus and just want to implement basic file reading or writing you should not start by reading that.

1.3.1 How do I write a NeXus file?

The NeXus Application Program Interface (API) provides a set of subroutines that make it easy to read and write NeXus files. These subroutines are available in C, Fortran 77, Fortran 90, Java, Python, C++, and IDL.

The API uses a very simple *state* model to navigate through a NeXus file. (Compare this example with NAPI Simple 2-D Write Example (C, F77, F90), in the NAPI chapter of Volume II, using the native HDF5 commands.) When you open a file, the API provides a file *handle*, which stores the current location, i.e. which group and/or field is currently open. Read and write operations then act on the currently open entity. Following the simple example of Simple example data file, we walk through a schematic of NeXus program written in C (without any error checking or real data).

Example 1.5 Writing a simple NeXus file using NAPI

```
#include "napi.h"

 int main()
 {
    NXhandle fileID;
    NXopen ("NXfile.nxs", NXACC_CREATE, &fileID);❶[]
      NXmakegroup (fileID, "Scan", "NXentry");❷[]
      NXopengroup (fileID, "Scan", "NXentry");
        NXmakegroup (fileID, "data", "NXdata");❸
        NXopengroup (fileID, "data", "NXdata");
        /* somehow, we already have arrays tth and counts, each length n*/
          NXmakedata (fileID, "two_theta", NX_FLOAT32, 1, &n);❹
          NXopendata (fileID, "two_theta");
            NXputdata (fileID, tth);❺
            NXputattr (fileID, "units", "degrees", 7, NX_CHAR);❻
          NXclosedata (fileID);  /* two_theta */❼
          NXmakedata (fileID, "counts", NX_FLOAT32, 1, &n);❽
          NXopendata (fileID, "counts");
            NXputdata (fileID, counts);
          NXclosedata (fileID);  /* counts */❾
        NXclosegroup (fileID);  /* data */
      NXclosegroup (fileID);  /* Scan */
    NXclose (&fileID);
    return;
}
```

❶ [line 6] Open the file `NXfile.nxs` with *create* access (implying write access). NAPI[4] returns a file identifier of type `NXhandle`.

❷ [line 7] Next, we create the NXentry group to contain the scan using `NXmakegroup()` and then open it for access using `NXopengroup()`. [5]

❸ [line 9] The plottable data is contained within an NXdata group, which must also be created and opened.

❹ [line 12] To create a field, call `NXmakedata()`, specifying the data name, type (`NX_FLOAT32`), rank (in this case, `1`), and length of the array (`n`). Then, it can be opened for writing. [6]

[3] http://svn.nexusformat.org/code/trunk/doc/api/NeXusIntern.pdf
[4] NAPI: The NeXus Application Programming Interface
[5] See the chapter about NeXus Base Classes for more information.
[6] The NeXus Data Types section describes the available data types, such as `NX_FLOAT32` and `NX_CHAR`.

5. [line 14] Write the data using `NXputdata()`.
6. [line 15] With the field still open, we can also add some data attributes, such as the data units[7], [8] which are specified as a character string (`type="NX_CHAR"`[9]) that is 7 bytes long.
7. [line 16] Then we close the field before opening another. In fact, the API will do this automatically if you attempt to open another field, but it is better style to close it yourself.
8. [line 17] The remaining fields in this group are added in a similar fashion. Note that the indentation whenever a new field or group are opened is just intended to make the structure of the NeXus file more transparent.
9. [line 20] Finally, close the groups (`NXdata` and `NXentry`) before closing the file itself.

1.3.2 How do I read a NeXus file?

Reading a NeXus file works in the same way by traversing the tree with the handle.

This schematic C code will read the two-theta array created in Writing a simple NeXus file using NAPI above. (Again, compare this example with one in the NAPI chapter of Volume II[10] using the native HDF5 commands.)

Example 1.6 Reading a simple NeXus file using NAPI

```
1  NXopen ('NXfile.nxs', NXACC_READ, &fileID);
2    NXopengroup (fileID, "Scan", "NXentry");
3      NXopengroup (fileID, "data", "NXdata");
4        NXopendata (fileID, "two_theta");
5          NXgetinfo (fileID, &rank, dims, &datatype);
6          NXmalloc ((void **) &tth, rank, dims, datatype);
7          NXgetdata (fileID, tth);
8        NXclosedata (fileID);
9      NXclosegroup (fileID);
10   NXclosegroup (fileID);
11 NXclose (fileID);
```

1.3.3 How do I browse a NeXus file?

NeXus files can also be viewed by a command-line browser, `nxbrowse`, which is included as a helper tool in the NeXus API distribution. The following is an example session of using `nxbrowse` to view a data file. The following commands are used in Example 1.7 in this session:

1. [line 1] Start `nxbrowse` from the UNIX command line and open file `lrcs3701.nxs` from IPNS/LRMECS.
2. [line 8] List the contents of the current group.
3. [line 11] Open the NeXus group `Histogram1`.
4. [line 23] Print the contents of the NeXus data labelled `title`.
5. [line 41] Close the current group.
6. [line 43] Quits `nxbrowse`.

[7] NeXus Data Units
[8] The NeXus rule about data units is described in the NeXus Data Units section.

[9] NXDL: categories of data types [??]
[10] Reading a simple NeXus file using native HDF5 commands

Example 1.7 Using nxbrowse

```
%> nxbrowse lrcs3701.nxs❶

NXBrowse 3.0.0. Copyright (C) 2000 R. Osborn, M. Koennecke, P. Klosowski
    NeXus_version = 1.3.3
    file_name = lrcs3701.nxs
    file_time = 2001-02-11 00:02:35-0600
    user = EAG/RO
NX> dir❷
  NX Group : Histogram1 (NXentry)
  NX Group : Histogram2 (NXentry)
NX> open Histogram1❸
NX/Histogram1> dir
  NX Data  : title[44] (NX_CHAR)
  NX Data  : analysis[7] (NX_CHAR)
  NX Data  : start_time[24] (NX_CHAR)
  NX Data  : end_time[24] (NX_CHAR)
  NX Data  : run_number (NX_INT32)
  NX Group : sample (NXsample)
  NX Group : LRMECS (NXinstrument)
  NX Group : monitor1 (NXmonitor)
  NX Group : monitor2 (NXmonitor)
  NX Group : data (NXdata)
NX/Histogram1> read title❹
  title[44] (NX_CHAR) = MgB2 PDOS 43.37g 8K 120meV E0@240Hz T0@120Hz
NX/Histogram1> open data
NX/Histogram1/data> dir
  NX Data  : title[44] (NX_CHAR)
  NX Data  : data[148,750] (NX_INT32)
  NX Data  : time_of_flight[751] (NX_FLOAT32)
  NX Data  : polar_angle[148] (NX_FLOAT32)
NX/Histogram1/data> read time_of_flight
  time_of_flight[751] (NX_FLOAT32) = [ 1900.000000 1902.000000 1904.000000 ...]
    units = microseconds
    long_name = Time-of-Flight [microseconds]
NX/Histogram1/data> read data
  data[148,750] (NX_INT32) = [ 1 1 0 ...]
    units = counts
    signal = 1
    long_name = Neutron Counts
    axes = polar_angle:time_of_flight
NX/Histogram1/data> close❺
NX/Histogram1> close
NX> quit❻
```

The source code of `nxbrowse`[11] provides an example of how to write a NeXus reader. The test programs included in the NeXus API may also be useful to study.

[11] https://svn.nexusformat.org/code/trunk/applications/NXbrowse/NXbrowse.c

Chapter 2

NeXus Design

This chapter actually defines the rules to use for writing valid NeXus files. An explanation of NeXus objects is followed by the definition of NeXus coordinate systems, the rules for structuring files and the rules for storing single items of data.

The structure of NeXus files is extremely flexible, allowing the storage both of simple data sets, such as a single data array and its axes, and also of highly complex data, such as the simulation results or an entire multi-component instrument. This flexibility is a necessity as NeXus strives to capture data from a wild variety of applications in x-ray, muSR and neutron scattering. The flexibility is achieved through a hierarchical structure, with related *fields*[1] collected together into *groups*, making NeXus files easy to navigate, even without any documentation. NeXus files are self-describing, and should be easy to understand, at least by those familiar with the experimental technique.

2.1 NeXus Objects and Terms

Before discussing the design of NeXus in greater detail it is necessary to define the objects and terms used by NeXus. These are:

Data Groups Group data fields and other groups together. Groups represent levels in the NeXus hierarchy

Data Fields Multidimensional arrays and scalars representing the actual data to be stored

Data Attributes Additional metadata which can be assigned to groups or data fields

Links Elements which point to data stored in another place in the file hierarchy

NeXus Base Classes Dictionaries of names possible in the various types of NeXus groups

NeXus Application Definitions Describe the minimum content of a NeXus file for a particular usage case

In the following sections these elements of NeXus files will be defined in more detail.

2.1.1 Data Groups

NeXus files consist of data groups, which contain fields and/or other groups to form a hierarchical structure. This hierarchy is designed to make it easy to navigate a NeXus file by storing related fields together. Data groups are identified both by a name, which must be unique within a particular group, and a class. There can be multiple groups with the same class but they must have different names (based on the HDF rules). For the class names used with NeXus data groups the prefix NX is reserved. Thus all NeXus class names start with NX.

[1] In this manual, we use the terms *field*, *data field*, and *data item* synonymously to be consistent with their meaning between NeXus data file instances and NXDL specification files.

2.1.2 Data Fields

Data fields contain the essential information stored in a NeXus file. They can be scalar values or multidimensional arrays of a variety of sizes (1-byte, 2-byte, 4-byte, 8-byte) and types (integers, floats, characters). The fields may store both experimental results (counts, detector angles, etc), and other information associated with the experiment (start and end times, user names, etc). Data fields are identified by their names, which must be unique within the group in which they are stored.

2.1.3 Data Attributes

Attributes are extra (meta-)information that are associated with particular fields. They are used to annotate the data, e.g. with physical units or calibration offsets, and may be scalar numbers or character strings. In addition, NeXus uses attributes to identify plottable data and their axes, etc. A description of possible attributes can be found in table data attributes. Finally, NeXus files themselves have global attributes which are listed in the global attributes table. that identify the NeXus version, file creation time, etc. Attributes are identified by their names, which must be unique in each field.

Table 2.1: NeXus File Global Attributes

Name	Type	Description
file_name	NX_CHAR	File name of original NeXus file to assist in identification if the external name has been changed
file_time	ISO 8601	Date and time of file creation
file_update_time	ISO 8601	Date and time of last file change at close
NeXus_version	NX_CHAR	Version of NeXus API used in writing the file
creator	NX_CHAR	Facility or program where the file originated

Table 2.2: NeXus Data Attributes. For the full specification of attributes, see the NXDL section in Volume II of the manual.

Name	Type	Description
units	NX_CHAR	Data units, given as character strings, must conform to the NeXus units standard. See the NeXus Data Units section for details.
signal	NX_INT	Defines which data set contains the signal to be plotted use `signal="1"` for main signal
axes	NX_CHAR	Defines the names of the dimension scales for this data set as a colon-delimited list. Note that some legacy data files may use a comma as delimiter. For example, suppose `data` is an array with elements `data[j][i]` (C) or `data(i,j)` (Fortran), with dimension scales `time_of_flight[i]` and `polar_angle[j]`, then `data` would have an attribute `axes="polar_angle:time_of_flight"` in addition to an attribute `signal="1"`.
axis	NX_INT	The original way of designating data for plotting, now superceded by the `axes` attribute. This defines the rank of the signal data for which this data set is a dimension scale in order of the fastest varying index (see a longer discussion in the section on `NXdata` structure), i.e. if the array being stored is `data`, with elements `data[j][i]` in C and `data(i,j)` in Fortran, `axis` would have the following values: ith dimension (`axis="1"`), jth dimension (`axis="2"`), etc.
primary	NX_INT32	Defines the order of preference for dimension scales which apply to the same rank of signal data. Use `primary="1"`to indicate preferred dimension scale
long_name	NX_CHAR	Defines title of signal data or axis label of dimension scale
calibration_status	NX_CHAR	Defines status of data value - set to "Nominal" or "Measured"
offset	NX_INT	Rank values off offsets to use for each dimension if the data is not in C storage order
stride	NX_INT	Rank values of steps to use when incrementing the dimension
vector	NX_FLOAT	3 values describing the axis of rotation or the direction of translation

Table 2.2: (continued)

Name	Type	Description
interpretation	NX_CHAR	Describes how to display the data. Allowed values include: `scaler` (0-D data), `spectrum` (1-D data), `image` (2-D data), or `vertex` (3-D data).

2.1.4 Links

Links are pointers to existing data somewhere else. The concept is very much like symbolic links in a unix filesystem. The NeXus definition sometimes requires to have access to the same data in different groups in the same file. For example: detector data is stored in the `NXinstrument/NXdetector` group but may be needed in `NXdata` for automatic plotting. Rather then replicating the data, NeXus uses links in such situations. See the figure for a more descriptive representation of the concept of linking.

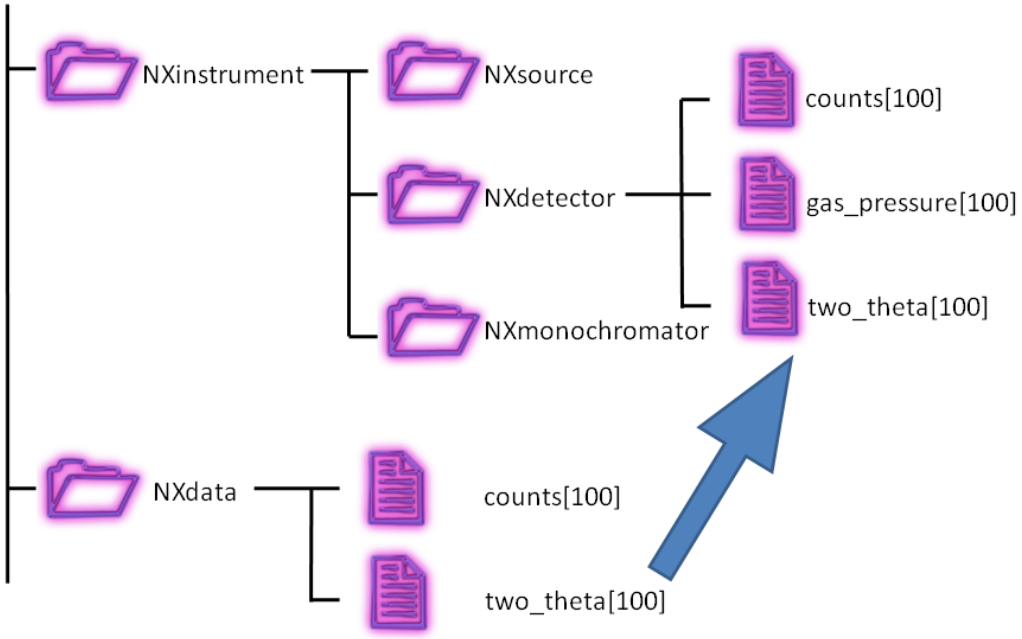

Figure 2.1: Linking in a NeXus file

2.1.5 NeXus Base Classes

Data groups often describe objects in the experiment (monitors, detectors, monochromators, etc.), so that the contents (both data fields and/or other data groups) comprise the properties of that object. NeXus has defined a set of standard objects, or base classes, out of which a NeXus file can be constructed. This is each data group is identified by a name and a class. The group class, defines the type of object and the properties that it can contain, whereas the group name defines a unique instance of that class. These classes are defined in XML using the NeXus Definition Language (NXDL) format. All NeXus class types adopted by the NIAC *must* begin with `NX`. Classes not adopted by the NIAC *must not* start with `NX`.

Note
NeXus base classes are the components used to build the NeXus data structure.

Not all classes define physical objects. Some refer to logical groupings of experimental information, such as plottable data, sample environment logs, beam profiles, etc. There can be multiple instances of each class. On the other hand, a typical NeXus file will only contain a small subset of the possible classes.

NeXus base classes are not proper classes in the same sense as used in object oriented programming languages. In fact the use of the term classes is actually misleading but has established itself during the development of NeXus. NeXus base classes are rather dictionaries of field names and their meanings which are permitted in a particular NeXus group implementing the NeXus class. This sounds complicated but becomes easy if you consider that most NeXus groups describe instrument components. Then for example, a NXmonochromator base class describes all the possible field names which NeXus allows to be used to describe a monochromator.

Most NeXus base classes represent instrument components. Some are used as containers to structure information in a file (`NXentry`, `NXcollection`, `NXinstrument`, `NXprocess`, `NXparameter`). But there are some base classes which have special uses which need to be mentioned here:

NXdata `NXdata` is used to identify the default plottable data. The notion of a default plot of data is a basic motivation of NeXus.

NXlog `NXlog` is used to store time stamped data like the log of a temperature controller. Basically you give a start time, and arrays with a difference in seconds to the start time and the values read.

NXnote This group provides a place to store general notes, images, video or whatever. A mime type is stored together with a binary blob of data. Please use this only for auxiliary information, for example an image of your sample, or a photo of your boss.

NXgeometry `NXgeometry` and its subgroups `NXtranslation`,`NXorientation`, `NXshape` are used to store absolute positions in the laboratory coordinate system or to define shapes.

These groups can appear anywhere in the NeXus hierarchy, where needed. Preferably close to the component they annotate or in a `NXcollection`. All of the base classes are documented in the reference manual.

2.1.5.1 `NXdata` Facilitates Automatic Plotting

The most notable special base class (or *group* in NeXus is `NXdata`. `NXdata` is the answer to a basic motivation of NeXus to facilitate automatic plotting of data. `NXdata` is designed to contain the main dataset and its associated dimension scales (axes) of a NeXus data file. The usage scenario is that an automatic data plotting program just opens a `NXentry` and then continues to search for any `NXdata` groups. These `NXdata` groups represent the plottable data. Here is the way an automatic plotting program ought to work:

1. Search for `NXentry` groups
2. Open an `NXentry`
3. Search for `NXdata` groups
4. Open an `NXdata` group
5. Identify the plottable data.
   (a) Search for a dataset with attribute `signal=1`. This is your main dataset. (There should be *only one* dataset that matches.)
   (b) Try to read the `axes` attribute of the main dataset, if it exists.
      i. The value of `axes` is a colon- or comma-separated list of the datasets describing the dimension scales (such as `axes="polar_angle:time_of_flight"`).
      ii. Parse `axes` and open the datasets to describe your dimension scales
   (c) If `axes` does not exist:
      i. Search for datasets with attributes `axis=1`, `axis=2`, etc. These are the datasets describing your axis. There may be several datasets for any axis, i.e. there may be multiple datasets with the attribute `axis=1`. Among them the dataset with the attribute `primary=1` is the preferred one. All others are alternative dimension scales.

ii. Open the datasets to describe your dimension scales.

6. Having found the default plottable data and its dimension scales: make the plot

2.1.6 NeXus Application Definitions

The objects described so far provide us with the means to store data from a wide variety of instruments, simulations or processed data as resulting from data analysis. But NeXus strives to express strict standards for certain applications of NeXus too. The tool which NeXus uses for the expression of such strict standards is the NeXus Application Definition. A NeXus Application Definition describes which groups and data items have to be present in a file in order to properly describe an application of NeXus. For example for describing a powder diffraction experiment. Typically an application definition will contain only a small subset of the many groups and fields defined in NeXus. NeXus application definitions are also expressed in the NeXus Definition Language (NXDL). A tool exists which allows one to validate a NeXus file against a given application definition.

Note
NeXus application definitions define the *minimum* information necessary to satisfy data analysis or other data processing.

Another way to look at a NeXus application definition is as a contract between a file producer (writer) and a file consumer (reader). The contract reads: *If you write your files following a particular NeXus application definition, I can process these files with my software.*

Yet another way to look at a NeXus application definition is to understand it as an interface definition between data files and the software which uses this file. Much like an interface in the Java or other modern object oriented programming languages.

In contrast to NeXus base classes, NeXus supports inheritance in application definitions.

Please note that a NeXus Application Definition will only define the bare minimum of data necessary to perform common analysis with data. Practical files will nearly always contain more data. One of the beauties of NeXus is that it is always possible to add more data to a file without breaking its compliance with its application definition.

2.2 NeXus Coordinate Systems

NeXus uses the *McStas coordinate system* as its laboratory coordinate system.

Coordinate systems in NeXus have undergone significant development. Initially, just motor positions of relevant motors were stored without further standardization. This soon proved to be to little and the *NeXus polar coordinate* system was developed. This system still is very close to angles meaningful to an instrument scientist but allows to define general positions of components easily. Then users from the simulation community approached the NeXus team and asked for a means to store absolute coordinates. This was implemented through the use of the *NXgeometry* class on top of the *McStas* system. We soon learned that all the things we do can be expressed through the McStas coordinate system. So it became the reference coordinate system for NeXus. `NXgeometry` was expanded to allow the description of shapes when the demand came up. Later members of the CIF team convinced the NeXus team of the beauty of transformation matrices and NeXus was enhanced to store the necessary information to fully map CIF concepts. Not much had to be changed though as we choose to document the existing angles in CIF terms. The CIF system allows to store arbitrary operations and nevertheless calculate absolute coordinates in the laboratory coordinate system. It also allows to convert from local, for example detector coordinate systems, to absolute coordinates in the laboratory system.

2.2.1 McStas and `NXgeometry` System

As stated above, NeXus uses the *McStas coordinate system* as its laboratory coordinate system. The instrument is given a global, absolute coordinate system where the z axis points in the direction of the incident beam, the x axis is perpendicular to the beam in the horizontal plane pointing left as seen from the source, and the y axis points upwards. See below for a drawing of the McStas coordinate system. The origin of this coordinate system is the sample position or, if this is ambiguous, the center

of the sample holder with all angles and translations set to zero. The McStas coordinate system is illustrated in figure McStas Coordinate System.

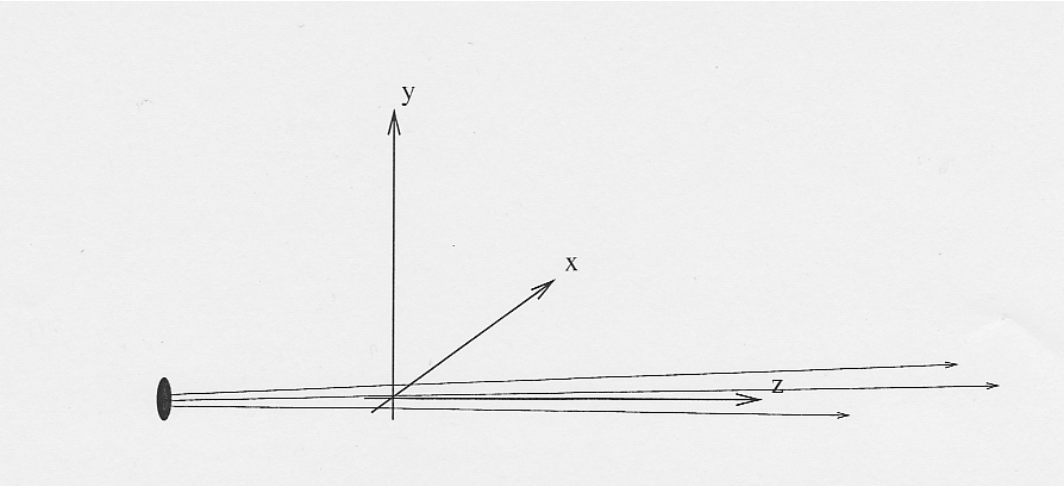

Figure 2.2: The McStas Coordinate System

Note
The NeXus definition of `+z` is opposite to that in the International Tables for Crystallography, volume G, and consequently, `+x` is also reversed.

The NeXus NXgeometry class directly uses the McStas coordinate system. NXgeometry classes can appear in any component in order to specify its position. The suggested name to use is geometry. In NXgeometry the NXtranslation/values field defines the absolute position of the component in the McStas coordinate system. The NXorientation/value field describes the orientation of the component as a vector of in the McStas coordinate system.

2.2.2 Simple (Spherical Polar) Coordinate System

In this system, the instrument is considered as a set of components through which the incident beam passes. The variable **distance** is assigned to each component and represents the effective beam flight path length between this component and the sample. A sign convention is used where negative numbers represent components pre-sample and positive numbers components post-sample. At each component there is local spherical coordinate system with the angles *polar_angle* and *azimuthal_angle*. The size of the sphere is the distance to the previous component.

In order to understand this spherical polar coordinate system it is helpful to look initially at the common condition that *azimuthal_angle* is zero. This corresponds to working directly in the horizontal scattering plane of the instrument. In this case *polar_angle* maps directly to the setting commonly known as *two theta*. Now, there are instruments where components live outside of the scattering plane. Most notably detectors. In order to describe such components we first apply the tilt out of the horizontal scattering plane as the *azimuthal_angle*. Then, in this tilted plane, we rotate to the component. The beauty of this is that *polar_angle* is always *two theta*. Which, in the case of a component out of the horizontal scattering plane, is not identical to the value read from the motor responsible for rotating the component. This situation is shown in Figure: Polar Coordinate System.

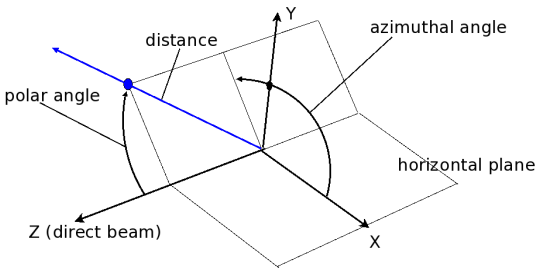

Figure 2.3: NeXus Simple (Spherical Polar) Coordinate System

2.2.3 Coordinate Transformations

Another way to look at coordinates is through the use of transformation matrices. In this world view, the absolute position of a component or a detector pixel with respect to the laboratory corrdinate system is calculated by applying a series of translations and rotations. These operations are commonly expressed as transformation matrices and their combination as matrix multiplication. A very important aspect is that the order of application of the individual operations *does* matter. Another important aspect is that any operation transforms the whole coordinate system and gives rise to a new local coordinate system. The mathematics behind this is well known and used in such applications such as industrial robot control, space flight and computer games. The beauty in this comes from the fact that the operations to apply map easily to instrument settings and constants. It is also easy to analyze the contribution of each individual operation: this can be studied under the condition that all other operations are at a zero setting.

In order to use coordinate transformations, several morsels of information need to be known:

Type The type of operation: rotation or translation

Direction The direction of the translation or the direction of the rotation axis

Value The angle of rotation or the length of the translation

Order The order of operations to apply to move a component into its place.

How NeXus describes the order of operations to apply has not yet been decided. The authors favourite scheme is to use a special field at each instrument component, named *transform* which describes the operations to apply to get the component into its position as a list of colon separated paths to the operations to apply relative to the current NXentry. For paths in the same group, only the name need to be given. Detectors may need two such fields: the transform field to get the detector as a whole into its position and a *transform_pixel* field which describes how the absolute position of a detector pixel can be calculated.

For the NeXus spherical coordinate system, the order is implicit and is given in the next example.

Example 2.1 implicit order of NeXus spherical coordinate system

```
azimuthal_angle:polar_angle:distance
```

This is also a nice example of the application of transformation matrices:

1. You first apply `azimuthal_angle` as a rotation around z. This rotates the whole coordinate out of the plane.
2. Then you apply `polar_angle` as a rotation around y in the tilted coordinate system.
3. This also moves the direction of the z vector. Along which you translate the component to place by distance.

2.3 Rules for Structuring Information in NeXus Files

All NeXus files contain one or many groups of type `NXentry` at root level. Many files contain only one `NXentry` group, then the name is `entry`. The NXentry level of hierarchy is there to support the storage of multiple related experiments in one file. Or to allow the NeXus file to serve as a container for storing a whole scientific workflow from data acquisition to publication ready data. Also, `NXentry` class groups can contain raw data or processed data. For files with more than one `NXentry` group, since HDF requires that no two items at the same level in an HDF file may have the same name, the NeXus fashion is to assign names with an incrementing index appended, such as `entry1`, `entry2`, `entry3`, etc.

In order to illustrate what is written in the text, example hierarchies like the one in figure Raw Data are provided.

2.3.1 Content of a Raw Data `NXentry` Group

An example raw data hierarchy is shown in figure Raw Data (only showing the relevant parts of the data hierarchy). In the example shown, the `data` field in the `NXdata` group is linked to the 2-D detector data (a 512x512 array of 32-bit integers) which has the attribute `signal=1`. Note that `[,]` represents a 2D array.

```
entry:NXentry
  instrument:NXinstrument
    source:NXsource
    ....
    detector:NXdetector
      data:NX_INT32[512,512]
        @signal = 1
  sample:NXsample
  control:NXmonitor
  data:NXdata
    data --> /entry/instrument/detector/data
```

Figure 2.4: NeXus Raw Data Hierarchy

An `NXentry` describing raw data contains at least a `NXsample`, one `NXmonitor`, one `NXdata` and a `NXinstrument` group. It is good practice to use the names `sample` for the `NXsample` group, `control` for the `NXmonitor` group holding the experiment controlling monitor and `instrument` for the `NXinstrument` group. The `NXinstrument` group contains further groups describing the individual components of the instrument as appropriate.

The `NXdata` group contains links to all those data items in the `NXentry` hierarchy which are required to put up a default plot of the data. As an example consider a SAXS instrument with a 2D detector. The `NXdata` will then hold a link to the detector image. If there is only one `NXdata` group, it is good practice to name it `data`. Otherwise, the name of the detector bank represented is a good selection.

2.3.2 Content of a processed data `NXentry` group

Processed data, see figure Processed Data, in this context means the results of a data reduction or data analysis program. Note that `[]` represents a 1D array.

```
entry:NXentry
  reduction:NXprocess
    program_name = "pyDataProc2010"
    version = "1.0a"
    input:NXparameter
      filename = "sn2013287.nxs"
  sample:NXsample
  data:NXdata
    data
      @signal = 1
```

Figure 2.5: NeXus Processed Data Hierarchy

NeXus stores such data in a simplified `NXentry` structure. A processed data `NXentry` has at minimum a `NXsample`, a `NXdata` and a `NXprocess` group. Again the preferred name for the `NXsample` group is `sample`. In the case of processed data, the `NXdata` group holds the result of the processing together with the associated axis data. The `NXprocess` group holds the name and version of the program used for this processing step and further `NXparameter` groups. These groups ought to contain the parameters used for this data processing step in suitable detail so that the processing step can be reproduced.

Optionally a processed data `NXentry` can hold a `NXinstrument` group with further groups holding relevant information about the instrument. The preferred name is again `instrument`. Whereas for a raw data file, NeXus strives to capture as much data as possible, a `NXinstrument` group for processed data may contain a much-reduced subset.

2.3.3 `NXsubentry` or Multi-Method Data

Especially at synchrotron facilities, there are experiments which perform several different methods on the sample at the same time. For example, combine a powder diffraction experiment with XAS. This may happen in the same scan, so the data needs to be grouped together. A suitable `NXentry` would need to adhere to two different application definitions. This leads to name clashes which cannot be easily resolved. In order to solve this issue, the following scheme was implemented in NeXus:

- The complete beamline (all data) is stored in an appropriate hierarchy in an `NXentry`.
- The `NXentry` group contains further `NXsubentry` groups, one for each method. Each `NXsubentry` group is constructed like a `NXentry` group. It contains links to all those data items required to fulfill the application definition for the particular method it represents.

See figure NeXus Multi Method Hierarchy for an example hierarchy. Note that `[,]` represents a 2D array.

```
entry:NXentry
  user:NXuser
  sample:NXsample
  instrument:NXinstument
    SASdet:NXdetector
      data:[,]
        @signal = 1
    fluordet:NXdetector
      data:[,]
        @signal = 1
    large_area:NXdetector
      data:[,]
  SAS:NXsubentry
    definition = "NXsas"
    instrument:NXinstrument
      detector:NXdetector
        data --> /entry/instrument/SASdet/data
    data:NXdata
      data --> /entry/instrument/SASdet/data
  Fluo:NXsubentry
    definition = "NXFluo"
    instrument:NXinstrument
      detector --> /entry/instrument/fluordet/data
      detector2 --> /entry/instrument/large_area/data
    data:NXdata
      detector --> /entry/instrument/fluordet/data
```

Figure 2.6: NeXus Multi Method Hierarchy

2.3.4 Rules for Special Cases

2.3.4.1 Scans

Scans are difficult to capture because they have great variety. Basically, any variable can be scanned. Such behaviour cannot be captured in application definitions. Therefore NeXus solves this difficulty with a set of rules. In this section, `NP` is used as a symbol for the number of scan points.

- The scan dimension `NP` is always the first dimension of any multi-dimensional dataset. The reason for this is that HDF allows the first dimension of a dataset to be unlimited. Which means, that data can be appended to the dataset during the scan.
- All data is stored as arrays of dimensions `NP`, original dimensions of the data at the appropriate position in the `NXentry` hierarchy.
- The `NXdata` group has to contain links to all variables varied during the scan and the detector data. Thus the `NXdata` group mimics the usual tabular representation of a scan.
- Datasets in an `NXdata` group must contain the proper attributes to enable the default plotting, as described in the section titled NXdata Facilitates Automatic Plotting.

2.3.4.1.1 Simple scan

Examples may be in order here. Let us start with a simple case, the sample is rotated around its rotation axis and data is collected in a single point detector. See figure Simple Scan for an overview. Then we have:

- A dataset at `NXentry/NXinstrument/NXdetector/data` of length `NP` containing the count data.
- A dataset at `NXentry/NXsample/rotation_angle` of length `NP` containing the positions of `rotation_angle` at the various steps of the scan.

- NXdata contains links to NXentry/NXinstrument/NXdetector/data and NXentry/NXsample/rotation_angle.
- All other data fields have their normal dimensions.

```
entry:NXentry
  instrument:NXinstrument
    detector:NXdetector
      data[NP]
        @signal = 1
  sample:NXsample
    rotation_angle[NP]
      @axis=1
  control:NXmonitor
    data[NP]
  data:NXdata
    data --> /entry/instrument/detector/data
    rotation_angle --> /entry/sample/rotation_angle
```

Figure 2.7: NeXus Simple Scan Example

2.3.4.1.2 Simple scan with area detector

The next example is the same scan but with an area detector with xsize times ysize pixels. The only thing which changes is that /NXentry/NXinstrument/NXdetector/data will have the dimensions NP, xsize, ysize. See figure Simple Scan with Area Detector for an overview.

```
entry:NXentry
  instrument:NXinstrument
    detector:NXdetector
      data:[NP,xsize,ysize]
        @signal = 1
  sample:NXsample
    rotation_angle[NP]
      @axis=1
  control:NXmonitor
    data[NP]
  data:NXdata
    data --> /entry/instrument/detector/data
    rotation_angle --> /entry/sample/rotation_angle
```

Figure 2.8: NeXus Simple Scan Example with Area Detector

2.3.4.1.3 Complex *hkl* scan

The next example involves a complex movement along an axis in reciprocal space which requires mutiple motors of a four circle diffractometer to be varied during the scan. We then have:

- A dataset at NXentry/NXinstrument/NXdetector/data of length NP containing the count data.
- A dataset at NXentry/NXinstrument/NXdetector/polar_angle of length NP containing the positions of the detector's polar_angle at the various steps of the scan.
- A dataset at NXentry/NXsample/rotation_angle of length NP containing the positions of rotation_angle at the various steps of the scan.

- A dataset at `NXentry/NXsample/chi` of length `NP` containing the positions of chi at the various steps of the scan.
- A dataset at `NXentry/NXsample/phi` of length `NP` containing the positions of phi at the various steps of the scan.
- A dataset at `NXentry/NXsample/h` of length `NP` containing the positions of the reciprocal coordinate `h` at the various steps of the scan.
- A dataset at `NXentry/NXsample/k` of length `NP` containing the positions of the reciprocal coordinate `k` at the various steps of the scan.
- A dataset at `NXentry/NXsample/l` of length `NP` containing the positions of the reciprocal coordinate `l` at the various steps of the scan.
- `NXdata` contains links to `NXentry/NXinstrument/NXdetector/data`, `NXentry/NXinstrument/NXdetector/polar_angle`, `NXentry/NXsample/rotation_angle`, `NXentry/NXsample/chi`, `NXentry/NXsample/phi`, `NXentry/NXsample/h`, `NXentry/NXsample/k` and `NXentry/NXsample/l` The datasets in `NXdata` must have the appropriate attributes as described in the axis location section.
- All other data fields have their normal dimensions.

```
entry:NXentry
  instrument:NXinstrument
    detector:NXdetector
      data[NP]
        @signal = 1
      polar_angle[NP]
        @axis = 1
      name
  sample:NXsample
    name
    rotation_angle[NP]
      @axis=1
    chi[NP]
      @axis=1
    phi[NP]
      @axis=1
    h[NP]
      @axis=1
      @primary=1
    k[NP]
      @axis=1
    l[NP]
      @axis=1
  control:NXmonitor
    data[NP]
  data:NXdata
    data --> /entry/instrument/detector/data
    rotation_angle --> /entry/sample/rotation_angle
    chi --> /entry/sample/chi
    phi --> /entry/sample/phi
    polar_angle --> /entry/instrument/detector/polar_angle
    h --> /entry/sample/h
    k --> /entry/sample/k
    l --> /entry/sample/l
```

Figure 2.9: NeXus Complex *hkl* Scan

2.3.4.1.4 Multi-parameter scan: XAS

Data can be stored almost anywhere in the NeXus tree. While the previous examples showed data arrays in either `NXdetector` or `NXsample`, this example demonstrates that data can be stored in other places. Links are used to reference the data.

The example is for X-ray Absorption Spectroscopy (XAS) data where the monochromator energy is step-scanned and counts are read back from detectors before (`I0`) and after (`I`) the sample. These energy scans are repeated at a sequence of sample temperatures to map out, for example, a phase transition. While it is customary in XAS to plot *log(I0/I)*, we show them separately here in two different `NXdata` groups to demonstrate that such things are possible. Note that the length of the 1-D energy array is `NE` while the length of the 1-D temperature array is `NT`

```
entry:NXentry
  instrument:NXinstrument
    I:NXdetector
      data:NX_NUMBER[NE,NT]
        @signal = 1
        @axes = "energy:temperature"
      energy --> /entry/monochromator/energy
      temperature --> /entry/sample/temperature
    I0:NXdetector
      data:NX_NUMBER[NE,NT]
        @signal = 1
        @axes = "energy:temperature"
      energy --> /entry/monochromator/energy
      temperature --> /entry/sample/temperature
  sample:NXsample
    temperature:NX_NUMBER[NT]
  monochromator:NXmonochromator
    energy:NX_NUMBER[NE]
  I_data:NXdata
    data --> /entry/instrument/I/data
    energy --> /entry/monochromator/energy
    temperature --> /entry/sample/temperature
  I0_data:NXdata
    data --> /entry/instrument/I00/data
    energy --> /entry/monochromator/energy
    temperature --> /entry/sample/temperature
```

Figure 2.10: NeXus Multi-parameter scan: XAS

2.3.4.2 Rastering

Rastering is the process of making experiments at various locations in the sample volume. Again, rasterisation experiments can be variable. Some people even raster on spirals! Rasterisation experiments are treated the same way as described above for scans. Just replace `NP` with `P`, the number of raster points.

Special rules apply if a rasterisation happens on a regular grid of size `xraster, yraster`. Then the variables varied in the rasterisation will be of dimensions `xraster, yraster` and the detector data of dimensions `xraster, yraster, (orginal dimensions)` of the detector. For example, an area detector of size `xsize, ysize` then it is stored with dimensions `xraster, yraster, xsize, ysize`.

Warning
Be warned: if you use the 2D rasterisation method with `xraster, yraster` you may end up with invalid data if the scan is aborted prematurely. This cannot happen if the first method is used.

2.3.4.3 NXcollection

On demand from the community, NeXus introduced a more informal method of storing information in a NeXus file. This is the `NXcollection` class which can appear anywhere underneath `NXentry`. `NXcollection` is a container for holding other data. The foreseen use is to document collections of similar data which do not otherwise fit easily into the `NXinstrument` or `NXsample` hierarchy, such as the intent to record *all* motor positions on a synchrotron beamline. Thus, `NXcollection` serves as a quick point of access to data for an instrument scientist or another expert. An example usage case is documented in figure NXcollection example.

```
entry:NXentry
  positioners:NXcollection
    mxx:NXpositioner
    mzz:NXpositioner
    sgu:NXpositioner
    ttv:NXpositioner
    hugo:NXpositioner
    ....
  scalars:NXcollection
    title (dataset)
    lieselotte (dataset)
    ...
  detectors:NXcollection
    Pilatus:NXdata
    MXX-45:NXdata
    ....
```

Figure 2.11: NXcollection Example

2.4 Rules for Storing Data Items in NeXus Files

This section describes the rules which apply for storing single data fields in data files.

2.4.1 Naming Conventions

Group and field Names used within NeXus follow a naming convention which is made up from the following rules: The names of NeXus *group*s and *field*s must only contain a restricted set of characters. This set may be described by this regular expression syntax regular expression syntax:

Example 2.2 Regular expression pattern for NXDL group and field names

```
[A-Za-z_][\w_]*
```

Note that this name pattern starts with a letter (upper or lower case) or "_" (underscore), then letters, numbers, and "_" and is limited to no more than 63 characters (imposed by the HDF5 rules for names).

Sometimes it is necessary to combine words in order to build a descriptive name for a data field or a group. In such cases lowercase words are connected by underscores.

```
number_of_lenses
```

For all data fields, only names from the NeXus base class dictionaries are to be used.[2] If a data field name or even a complete component is missing, please suggest the addition to the NIAC. The addition will usually be accepted provided it is not a duplication of an existing field and adequately documented.

[2]The NeXus base classes provide a comprehensive dictionary of terms than can be used for each class.

2.4.2 NeXus Array Storage Order

NeXus stores multi dimensional arrays of physical values in C language storage order, where the last dimension is the fastest varying. This is the rule. *Good reasons are required to deviate from this rule.*

It is possible to store data in storage orders other than C language order. As well it is possible to specify that the data needs to be converted first before being useful. Consider one situation, when data must be streamed to disk as fast as possible and conversion to C language storage order causes unnecessary latency. This case presents a good reason to make an exception to the standard rule.

2.4.2.1 Non C Storage Order

In order to indicate that the storage order is different from C storage order two additional data set attributes, offset and stride, have to be stored which together define the storage layout of the data. Offset and stride contain rank numbers according to the rank of the multidimensional data set. Offset describes the step to make when the dimension is multiplied by 1. Stride defines the step to make when incrementing the dimension. This is best explained by some examples.

Example 2.3 Offset and Stride for 1 D data:

```
* raw data = 0 1 2 3 4 5 6 7 8 9
   size[1] = { 10 }  // assume uniform overall array dimensions

* default stride:
   stride[1] = { 1 }
   offset[1] = { 0 }
   for i:
      result[i]:
         0 1 2 3 4 5 6 7 8 9

* reverse stride:
   stride[1] = { -1 }
   offset[1] = { 9 }
   for i:
      result[i]:
         9 8 7 6 5 4 3 2 1 0
```

Example 2.4 Offset and Stride for 2D Data

```
* raw data = 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19
   size[2] = { 4, 5 }  // assume uniform overall array dimensions

* row major (C) stride:
   stride[2] = { 5, 1 }
   offset[2] = { 0, 0 }
   for i:
      for j:
         result[i][j]:
            0 1 2 3 4
            5 6 7 8 9
            10 11 12 13 14
            15 16 17 18 19

* column major (Fortran) stride:
   stride[2] = { 1, 4 }
   offset[2] = { 0, 0 }
   for i:
      for j:
         result[i][j]:
            0 4 8 12 16
            1 5 9 13 17
            2 6 10 14 18
            3 7 11 15 19

* "crazy reverse" row major (C) stride:
   stride[2] = { -5, -1 }
   offset[2] = { 4, 5 }
   for i:
      for j:
         result[i][j]:
            19 18 17 16 15
            14 13 12 11 10
            9 8 7 6 5
            4 3 2 1 0
```

Example 2.5 Offset and Stride for 3D Data

```
 * raw data = 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19
        20 21 22 23 24 25 26 27 28 29 30 31 32 33 34 35 36 37 38 39
        40 41 42 43 44 45 46 47 48 49 50 51 52 53 54 55 56 57 58 59
     size[3] = { 3, 4, 5 }  // assume uniform overall array dimensions

  * row major (C) stride:
     stride[3] = { 20, 5, 1 }
     offset[3] = { 0, 0, 0 }
     for i:
        for j:
           for k:
              result[i][j][k]:
                 0 1 2 3 4
                 5 6 7 8 9
                 10 11 12 13 14
                 15 16 17 18 19

                 20 21 22 23 24
                 25 26 27 28 29
                 30 31 32 33 34
                 35 36 37 38 39

                 40 41 42 43 44
                 45 46 47 48 49
                 50 51 52 53 54
                 55 56 57 58 59

  * column major (Fortran) stride:
     stride[3] = { 1, 3, 12 }
     offset[3] = { 0, 0, 0 }
     for i:
        for j:
           for k:
              result[i][j][k]:
                 0 12 24 36 48
                 3 15 27 39 51
                 6 18 30 42 54
                 9 21 33 45 57

                 1 13 25 37 49
                 4 16 28 40 52
                 7 19 31 43 55
                 10 22 34 46 58

                 2 14 26 38 50
                 5 17 29 41 53
                 8 20 32 44 56
                 11 23 35 47 59
```

2.4.3 NeXus Data Types

Table 2.3: Matching regular expressions for NeXus data types

description	matching regular expression
integer	`NX_INT(8|16|32|64)`
floating-point	`NX_FLOAT(32|64)`
array	`(\[0-9\])?`

Table 2.3: (continued)

description	matching regular expression
valid item name	`^[A-Za-z_][A-Za-z0-9_]*$`
valid class name	`^NX[A-Za-z0-9_]*$`

NeXus supports numeric data as either integer or floating-point numbers. A number follows that indicates the number of bits in the word. The table above shows the regular expressions that matches the data type specifier.

integers

`NX_INT8`, `NX_INT16`, `NX_INT32`, or `NX_INT64`

floating-point numbers

`NX_FLOAT32` or `NX_FLOAT64`

date / time stamps

`NX_DATE_TIME` or `ISO8601`: Dates and times are specified using ISO-8601 standard definitions. Refer to Section 2.4.3.1.

strings

All strings are to be encoded in UTF-8. Since most strings in a NeXus file are restricted to a small set of characters and the first 128 characters are standard across encodings, the encoding of most of the strings in a NeXus file will be a moot point. Where encoding in UTF-8 will be important is when recording peoples names in `NXuser` and text notes in `NXnotes`.

Because the few places where encoding is important also have unpredictable content, as well as the way in which current operating systems handle character encoding, it is practically impossible to test the encoding used. Hence, `nxvalidate` provides no messages relating to character encoding.

binary data

Binary data is to be written as `UINT8`.

images

Binary image data is to be written using `UINT8`, the same as binary data, but with an accompanying image mime-type. If the data is text, the line terminator is `[CR][LF]`.

2.4.3.1 NeXus dates and times

NeXus dates and times should be stored using the ISO 8601[3] format e.g. `1996-07-31T21:15:22+0600`. The standard also allows for time intervals in fractional seconds with *1 or more digits of precision*. This avoids confusion, e.g. between U.S. and European conventions, and is appropriate for machine sorting.

Example 2.6 strftime() format specifiers for ISO-8601 time

```
%Y-%m-%dT%H:%M:%S%z
```

Note

Note that the `T` appears literally in the string, to indicate the beginning of the time element, as specified in ISO 8601. It is common to use a space in place of the `T`, such as `1996-07-31 21:15:22+0600`. While human-readable, compatibility with the ISO 8601 standard is not assured with this substitution. The `strftime()` format specifier for this is "`%Y-%m-%d %H:%M:%S%z`".

[3] ISO 8601: http://www.w3.org/TR/NOTE-datetime

2.4.4 NeXus Data Units

Given the plethora of possible applications of NeXus, it is difficult to define units to use. Therefore, the general rule is that you are free to store data in any unit you find fit. However, any data field must have a units attribute which describes the units, Wherever possible, SI units are preferred. NeXus units are written as a string attribute (`NX_CHAR`) and describe the engineering units. The string should be appropriate for the value. Values for the NeXus units must be specified in a format compatible with Unidata UDunits [4] Application definitions may specify units to be used for fields using an `enumeration`.

2.4.5 Linking Multi Dimensional Data with Axis Data

NeXus allows to store multi dimensional arrays of data. In most cases it is not sufficient to just have the indices into the array as a label for the dimensions of the data. Usually the information which physical value corresponds to an index into a dimension of the multi dimensional data set. To this purpose a means is needed to locate appropriate data arrays which describe what each dimension of a multi dimensional data set actually corresponds too. There is a standard HDF facility to do this: it is called dimension scales. Unfortunately, at a time, there was only one global namespace for dimension scales. Thus NeXus had to come up with its own scheme for locating axis data which is described here. A side effect of the NeXus scheme is that it is possible to have multiple mappings of a given dimension to physical data. For example a TOF data set can have the TOF dimension as raw TOF or as energy.

There are two methods of linking each data dimension to its respective dimension scale. The preferred method uses the `axes` attribute to specify the names of each dimension scale. The original method uses the `axis` attribute to identify with an integer the axis whose value is the number of the dimension. After describing each of these methods, the two methods will be compared. A prerequisite for both methods is that the data fields describing the axis are stored together with the multi dimensional data set whose axes need to be defined in the same NeXus group. If this leads to data duplication, use links.

2.4.5.1 Linking by name using the `axes` attribute

The preferred method is to define an attribute of the data itself called *axes*. The `axes` attribute contains the names of each dimension scale as a colon (or comma) separated list in the order they appear in C. For example:

Example 2.7 Preferred way of denoting axes

```
data:NXdata
  time_of_flight = 1500.0 1502.0 1504.0 ...
  polar_angle = 15.0 15.6 16.2 ...
  some_other_angle = 0.0 0.0 2.0 ...
  data = 5 7 14 ...
    @axes = polar_angle:time_of_flight
    @signal = 1
```

2.4.5.2 Linking by dimension number using the `axis` attribute

The original method is to define an attribute of each dimension scale called *axis*. It is an integer whose value is the number of the dimension, in order of fastest varying dimension. That is, if the array being stored is data with elements `data[j][i]` in C and `data(i,j)` in Fortran, where `i` is the time-of-flight index and `j` is the polar angle index, the `NXdata` group would contain:

[4] The UDunits specification also includes instructions for derived units. At present, the contents of NeXus `units` attributes are not validated in data files.

Example 2.8 Original way of denoting axes

```
data:NXdata
  time_of_flight = 1500.0 1502.0 1504.0 ...
    @axis = 1
    @primary = 1
  polar_angle = 15.0 15.6 16.2 ...
    @axis = 2
    @primary = 1
  some_other_angle = 0.0 0.0 2.0 ...
    @axis = 1
  data = 5 7 14 ...
    @signal = 1
```

The `axis` attribute must be defined for each dimension scale. The `primary` attribute is unique to this method of linking.

There are limited circumstances in which more than one dimension scale for the same data dimension can be included in the same `NXdata` group. The most common is when the dimension scales are the three components of an *(hkl)* scan. In order to handle this case, we have defined another attribute of type integer called `primary` whose value determines the order in which the scale is expected to be chosen for plotting, i.e.

1st choice: `primary="1"`
2nd choice: `primary="2"`
etc.

If there is more than one scale with the same value of the `axis` attribute, one of them must have set `primary="1"`. Defining the `primary` attribute for the other scales is optional.

Note
The `primary` attribute can only be used with the first method of defining dimension scales discussed above. In addition to the `signal` data, this group could contain a data set of the same rank and dimensions called `errors` containing the standard deviations of the data.

2.4.5.3 Discussion of the two linking methods

In general the method using the `axes` attribute on the multi dimensional data set should be preferred. This leaves the actual axis describing data sets unannotated and allows them to be used as an axis for other multi dimensional data. This is especially a concern as an axis describing a data set may be linked into another group where it may describe a completely different dimension of another data set.

Only when alternative axes definitions are needed, the `axis` method should be used to specify an axis of a data set. This is shown in the example above for the `some_other_angle` field where `axis="1"` denotes another possible primary axis for plotting. The default axis for plotting carries the `primary="1"` attribute.

Both methods of linking data axes will be supported in NeXus utilities that identify dimension scales, such as `NXUfindaxis()`.

2.4.6 Storing Detectors

There are very different types of detectors out there. Storing their data can be a challenge. As a general guide line: if the detector has some well defined form, this should be reflected in the data file. A linear detector becomes a linear array, a rectangular detector becomes an array of size `xsize` times `ysize`. Some detectors are so irregular that this does not work. Then the detector data is stored as a linear array, with the index being detector number till `ndet`. Such detectors must be accompanied by further arrays of length `ndet` which give `azimuthal_angle, polar_angle and distance` for each detector.

If data from a time of flight (TOF) instrument must be described, then the TOF dimension becomes the last dimension, for example an area detector of `xsize` *vs.* `ysize` is stored with TOF as an array with dimensions `xsize, ysize, ntof`.

2.4.7 Monitors are Special

Monitors , detectors that measure the properties of the experimental probe rather than the sample, have a special place in NeXus files. Monitors are crucial to normalize data. To emphasize their role, monitors are not stored in the `NXinstrument` hierarchy but on `NXentry` level in their own groups as there might be multiple monitors. Of special importance is the monitor in a group called `control`. This is the main monitor against which the data has to be normalized. This group also contains the counting control information, i.e. counting mode, times, etc.

Monitor data may be multidimensional. Good examples are scan monitors where a monitor value per scan point is expected or time-of-flight monitors.

2.4.8 Find the plottable data

Any program whose aim is to identify plottable data should use the following procedure:

1. Open the first top level NeXus group with class `NXentry`.
2. Open the first NeXus group with class `NXdata`.
3. Loop through NeXus fields in this group searching for the item with attribute `signal="1"` indicating this field has the plottable data.
4. Check to see if this field has an attribute called `axes`. If so, the attribute value contains a colon (or comma) delimited list (in the C-order of the data array) with the names of the dimension scales associated with the plottable data. And then you can skip the next two steps.
5. If the `axes` attribute is not defined, search for the one-dimensional NeXus fields with attribute `primary="1"`.
6. These are the dimension scales to label the axes of each dimension of the data.
7. Link each dimension scale to the respective data dimension by the `axis` attribute (`axis="1"`, `axis="2"`, ... up to the rank of the data).
8. If necessary, close the `NXdata` group, open the next one and repeat steps 3 to 6.
9. If necessary, close the `NXentry` group, open the next one and repeat steps 2 to 7.

Consult the NeXus API section, which describes the routines available to program these operations. In the course of time, generic NeXus browsers will provide this functionality automatically.

2.5 Physical File format

This section describes how NeXus structures are mapped to features of the underlying physical file format. This is a guide for people who wish to create NeXus files without using the NeXus-API.

2.5.1 Choice of HDF as Underlying File Format

At its beginnings, the founders of NeXus identified the Hierarchical Data Format (HDF), initially from the National Center for Supercomputing Applications (NCSA) at the University of Illinois at Urbana-Champaign (UIUC) and later spun off into its own group called The HDF Group (THG),[5] as a multi-platform data storage format with capacity for conveying large data payloads and a substantial user community. HDF (now HDF5) was provided with software to read and write data (this is the application-programmer interface, or API) using a large number of computing systems in common use for neutron and X-ray science. HDF is a binary data file format that supports compression and structured data.

[5] The HDF Group: http://www.hdfgroup.org/

2.5.2 Mapping NeXus into HDF

NeXus data structures map directly to HDF structures. NeXus *groups* are HDF4 *vgroups* or HDF5 *groups*, NeXus data sets (or *fields*) are HDF4 *SDS (scientific data sets)* or HDF5 *datasets*. Attributes map directly to HDF group or dataset attributes.

The only special case is the NeXus class name. HDF4 supports a group class which is set with the `Vsetclass()` call and read with `VGetclass()`. HDF-5 has no group class. Thus the NeXus class is stored as an attribute to the HDF-5 group with the name `NX_class` and value of the NeXus class name.

A NeXus `link` directly maps to the HDF linking mechanisms.

Note
Examples are provided in the Examples of reading or writing NeXus data files chapter of Volume II of this manual. These examples include software to write and read NeXus data files using the NAPI, as well as other software examples that use native (non-NAPI) libraries. In some cases the examples show the content of the NeXus data files that are produced. Here are links to some of the examples:

- Writing a simple NeXus file using NAPI
- File read example using NAPI
- Writing a simple NeXus file using native HDF5 commands
- Reading a simple NeXus file using native HDF5 commands
- Write a NeXus HDF5 file using Python with h5py
- Read a NeXus HDF5 file using Python with h5py

Perhaps the easiest way to view the implementation of NeXus in HDF5 is to view how the data structures look. For this, we use the `h5dump` command-line utility provided with the HDF5 support libraries. Short examples are provided for the basic NeXus data components:

- h5dump of a NeXus NXentry group: created in C NAPI by `NXmakegroup (fileID, "entry", "NXentry");`
- h5dump of a NeXus field (HDF5 dataset): created in C NAPI by `NXmakedata (fileID, "two_theta", NX_FLOAT32, 1, &n); NXopendata (fileID, "two_theta"); NXputdata (fileID, tth);`
- h5dump of a NeXus attribute: created in C NAPI by `NXputattr (fileID, "units", "degrees", 7, NX_CHAR);`
- h5dump of a NeXus link

See the sections NAPI Simple 2-D Write Example (C, F77, F90) and NAPI Simple 3-D Write Example (Python) in the Examples of reading or writing NeXus data files chapter of Volume II for examples that use the native HDF5 calls to write NeXus data files.

Example 2.9 `h5dump` of a NeXus `NXentry` group

```
GROUP "entry" {
  ATTRIBUTE "NX_class" {
   DATATYPE  H5T_STRING {
       STRSIZE 7;
       STRPAD H5T_STR_NULLPAD;
       CSET H5T_CSET_ASCII;
       CTYPE H5T_C_S1;
     }
   DATASPACE  SCALAR
   DATA {
   (0): "NXentry"
   }
  }
  # ... group contents
}
```

Example 2.10 `h5dump` of a NeXus field (HDF5 dataset)

```
 DATASET "two_theta" {
  DATATYPE  H5T_IEEE_F64LE
  DATASPACE  SIMPLE { ( 31 ) / ( 31 ) }
  DATA {
  (0): 17.9261, 17.9259, 17.9258, 17.9256, 17.9254, 17.9252,
  (6): 17.9251, 17.9249, 17.9247, 17.9246, 17.9244, 17.9243,
  (12): 17.9241, 17.9239, 17.9237, 17.9236, 17.9234, 17.9232,
  (18): 17.9231, 17.9229, 17.9228, 17.9226, 17.9224, 17.9222,
  (24): 17.9221, 17.9219, 17.9217, 17.9216, 17.9214, 17.9213,
  (30): 17.9211
  }
  ATTRIBUTE "units" {
     DATATYPE  H5T_STRING {
       STRSIZE 7;
       STRPAD H5T_STR_NULLPAD;
       CSET H5T_CSET_ASCII;
       CTYPE H5T_C_S1;
      }
     DATASPACE  SCALAR
     DATA {
     (0): "degrees"
     }
  }
  # ... other attributes
 }
```

Example 2.11 `h5dump` of a NeXus attribute

```
ATTRIBUTE "axes" {
   DATATYPE  H5T_STRING {
     STRSIZE 9;
     STRPAD H5T_STR_NULLPAD;
     CSET H5T_CSET_ASCII;
     CTYPE H5T_C_S1;
    }
   DATASPACE  SCALAR
   DATA {
   (0): "two_theta"
   }
}
```

Example 2.12 `h5dump` of a NeXus link

```
# NeXus links have two parts in HDF5 files.

# The dataset is created in some group.
# A "target" attribute is added to indicate the HDF5 path to this dataset.

ATTRIBUTE "target" {
   DATATYPE  H5T_STRING {
     STRSIZE 21;
     STRPAD H5T_STR_NULLPAD;
     CSET H5T_CSET_ASCII;
     CTYPE H5T_C_S1;
    }
   DATASPACE  SCALAR
   DATA {
   (0): "/entry/data/two_theta"
   }
}

# then, the hard link is created that refers to the original dataset
# (Since the name is "two_theta" in this example, it is understood that
# this link is created in a different HDF5 group than "/entry/data".)

DATASET "two_theta" {
   HARDLINK "/entry/data/two_theta"
}
```

2.5.3 Mapping NeXus into XML

This takes a bit more work than HDF. At the root of NeXus XML file is a XML element with the name `NXroot`. Further XML attributes to `NXroot` define the NeXus file level attributes. An example NeXus XML data file is provided in the Introduction chapter as Example Example 1.1

NeXus groups are encoded into XML as elements with the name of the NeXus class and an XML attribute `name` which defines the NeXus name of the group. Further group attributes become XML attributes. An example:

Example 2.13 NeXus group element in XML

```
  <NXentry name="entry">
  </NXentry>
```

NeXus data sets are encoded as XML elements with the name of the data. An attribute `NAPItype` defines the type and dimensions of the data. The actual data is stored as `PCDATA`[6] in the element. Another example:

Example 2.14 NeXus data elements

```
<mode NAPItype="NX_CHAR[7]">
  monitor
</mode>
<counts NAPItype="NX_INT32[4]">
  21 456  127876 319
</counts>
```

Data are printed in appropriate formats and in C storage order. The codes understood for `NAPItype` are all the NeXus data type names. The dimensions are given in square brackets as a comma separated list. No dimensions need to be given if the data is just a single value. Data attributes are represented as XML attributes. If the attribute is not a text string, then the attribute is given in the form: *type:value*, for example: `signal="NX_INT32:1"`.

NeXus links are stored in XML as XML elements with the name `NAPIlink` and a XML attribute `target` which stores the path to the linked entity in the file. If the item is linked under a different name, then this name is specified as a XML attribute name to the element `NAPIlink`.

The authors of the NeXus API worked with the author of the miniXML XML library to create a reasonably efficient way of handling numeric data with XML. Using the NeXus API handling something like 400 detectors versus 2000 time channels in XML is not a problem. But you may hit limits with XML as the file format when data becomes to large or you try to process NeXus XML files with general XML tools. General XML tools are normally ill prepared to process large amounts of numbers.

2.5.4 Special Attributes

NeXus makes use of some special attributes for its internal purposes. These attributes are stored as normal group or data set attributes in the respective file format. These are:

target This attribute is automatically created when items get linked. The target attribute contains a text string with the path to the source of the item linked.

napimount The `napimount` attribute is used to implement external linking in NeXus. The string is a URL to the file and group in the external file to link too. The system is meant to be extended. But as of now, the only format supported is: `nxfile://path-to-file#path-infile`. This is a NeXus file in the file system at path-to-file and the group path-infile in that NeXus file.

NAPIlink NeXus supports linking items in another group under another name. This is only supported natively in HDF-5. For HDF-4 and XML a crutch is needed. This crutch is a special class name or attribute `NAPIlink` combined with the target attribute. For groups, `NAPILink` is the group class, for data items a special attribute with the name `NAPIlink`.

[6] `PCDATA` is the XML term for *parsed character data* (see: http://www.w3schools.com/xml/xml_cdata.asp).

Chapter 3

Constructing NeXus Files and Application Definitions

In , we discussed the design of the NeXus format in general terms. In this section a more tutorial style introduction in how to construct a NeXus file is given. As an example a hypothetical instrument named WONI will be used.

Note

If you are looking for a tutorial on reading or writing NeXus data files using the NeXus API, consult the NAPI: The NeXus Application Programming Interface chapter of Volume II. For code examples, refer to Code Examples that use the NAPI chapter of Volume II. Alternatively, there are examples in the Example NeXus programs using native HDF5 commands chapter of writing and reading NeXus data files using the native HDF5 interfaces in C. Further, there are also some Python examples using the `h5py` package in the Complete h5py Example section.

3.1 The WOnderful New Instrument (WONI)

Consider yourself to be responsible for some hypothetical WOnderful New Instrument (WONI). You are tasked to ensure that WONI will record data according to the NeXus standard. For the sake of simplicity, WONI bears a strong resemblance to a simple powder diffractometer, but let's pretend that WONI cannot use any of the existing NXDL application definitions.

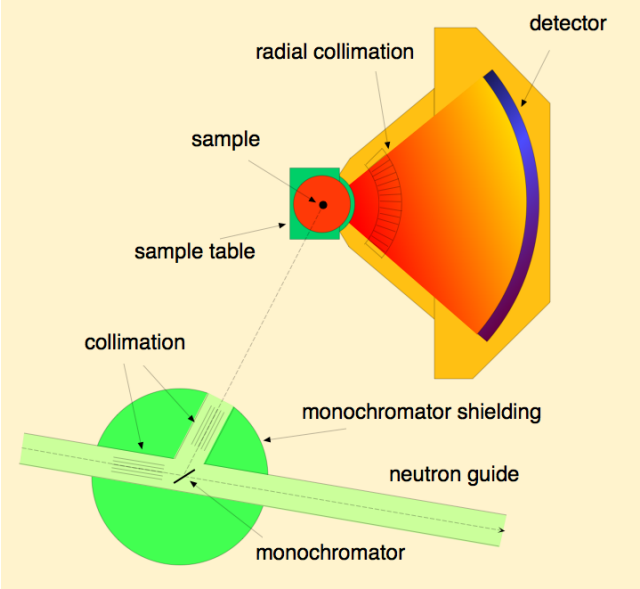

Figure 3.1: The (fictional) WONI example powder diffractometer

WONI uses collimators and a monochromator to illuminate the sample with neutrons of a selected wavelength as described in Figure 3.1. The diffracted beam is collected in a large, banana-shaped, position sensitive detector. Typical data looks like Figure 3.2. There is a generous background to the data plus quite a number of diffraction peaks.

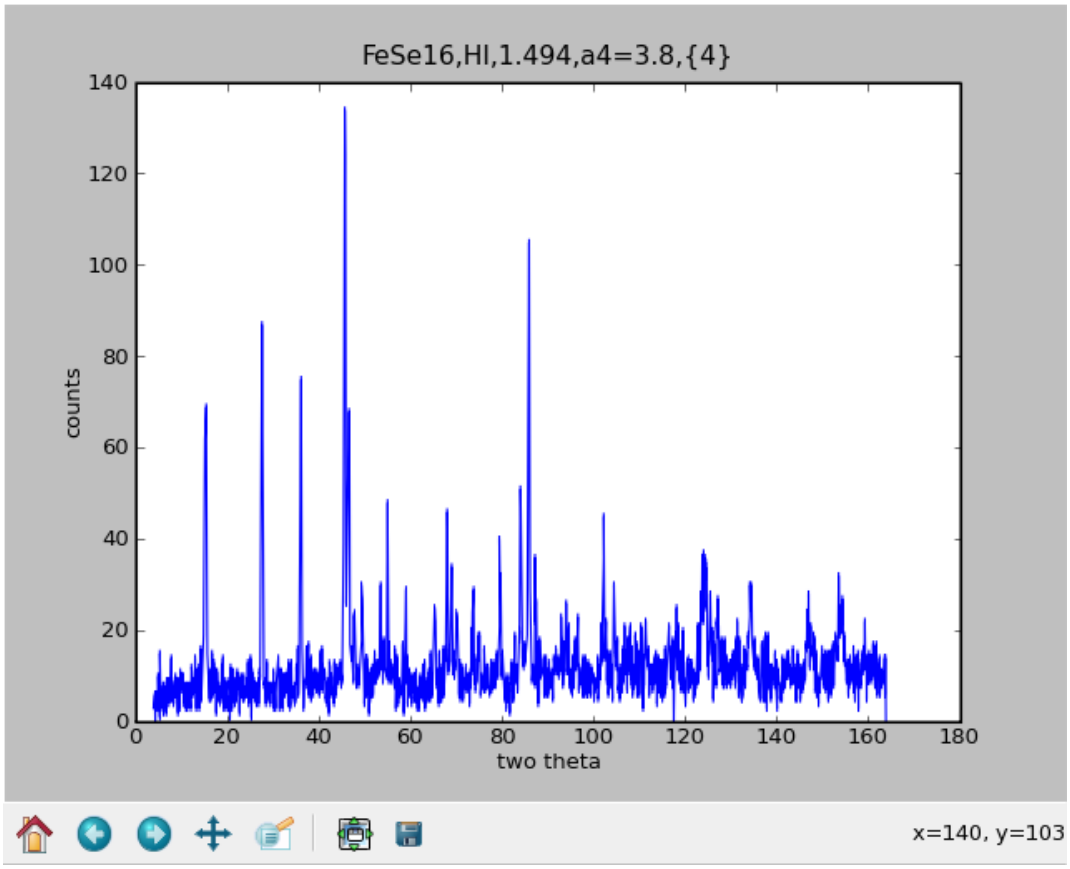

Figure 3.2: Example Powder Diffraction Plot from (fictional) WONI at HYNES

3.2 Constructing a NeXus file for WONI

The starting point for a NeXus file for WONI will be an empty basic NeXus file hierarchy as documented in figure Figure 3.3. In order to arrive at a full neXus file the following steps are required:

1. For each instrument component, decide which parameters need to be stored
2. Map the component parameters to NeXus groups and parameters and add the components to the `NXinstrument` hierarchy
3. Decide what needs to go into `NXdata`
4. Fill the `NXsample` and `NXmonitor` groups

```
        entry:NXentry
    NXdata
    NXinstrument
    NXmonitor
    NXsample

```

Figure 3.3: Basic structure of a NeXus file

3.2.1 Decide which parameters need to be stored

Now the various groups of this empty NeXus file shell need to be filled. The next step is to look at a design drawing of WONI. Identify all the instrument components like collimators, detectors, monochromators etc. For each component decide which values need to be stored. As NeXus aims to describe the experiment as good as possible, strive to capture as much information as practical.

3.2.2 Mapping parameters to NeXus

With the list of parameters to store for each component, consult the reference manual section on the NeXus base classes. You will find that for each of your instruments components there will be a suitable NeXus base class. Add this base class together with a name as a group under NXinstrument in your NeXus file hierarchy. Then consult the possible parameter names in the NeXus base class and match them with the parameters you wish to store for your instruments components.

As an example, consider the monochromator. You may wish to store: the wavelength, the d-value of the reflection used, the type of the monochromator and its angle towards the incoming beam. The reference manual tells you that NXcrystal is the right base class to use. Suitable fields for your parameters can be found in there to. After adding them to the basic NeXus file the file looks like in figure Figure 3.4

```
        entry:NXentry
    NXdata
    NXinstrument
       monochromator:Nxcrystal
          wavelength
          d_spacing
          rotation_angle
          reflection
          type
    NXmonitor
    NXsample

```

Figure 3.4: Basic structure of a NeXus file with a monochromator added

If a parameter or even a whole group is missing in order to describe your experiment, do not despair! Contact the NIAC and suggest to add the group or parameter. Give a little documentation what it is for. The NIAC will check that your suggestion is no duplicate and sufficiently documented and will then proceed to enhance the base classes with your suggestion.

A more elaborate example of the mapping process is given in the section Creating a NXDL Specification.

3.2.3 Decide on `NXdata`

The `NXdata/` group is supposed to contain the data required to put up a quick plot. For WONI this is a plot of counts versus two theta (polar_angle in NeXus) as can be seen in Figure 3.2. Now, in `NXdata`, create links to the appropriate data items in the `NXinstrument` hierarchy. In the case of WONI, both parameters live in the `detector:NXdetector` group.

3.2.4 Fill in auxiliary Information

Look at the section on `NXsample` in the NeXus reference manual. Choose appropriate parameters to store for your samples. Probably at least the name will be needed.

In order to normalize various experimental runs against each other it is necessary to know about the counting conditions and especially the monitor counts of the monitor used for normalization. The NeXus convention is to store such information in a `control:NXmonitor` group at `NXentry` level. Consult the reference for `NXmonitor` for field names. If additional monitors exist within your experiment, they will be stored as additional `NXmonitor` groups at entry level.

Consult the documentation for `NXentry` in order to find out under which names to store information such as titles, user names, experiment times etc.

A more elaborate example of this process can be found in the following section on creating an application definition.

3.3 Creating a NXDL Specification

An NXDL specification for a NeXus file is required if you desire to standardize NeXus files from various sources. Another name for a NXDL description is application definition. A NXDL specification can be used to verify NeXus files to conform to the standard encapsulated in the application definition. The process for constructing a NXDL specification is similar to the one described above for the construction of NeXus files.

One easy way to describe how to store data in the NeXus class structure and to create a NXDL specification is to work through an example. Along the way, we will describe some key decisions that influence our particular choices of metadata selection and data organization. So, on with the example ...

3.3.1 Application Definition Steps

With all this introductory stuff out of the way, let us look at the process required to define an application definition:

1. *Think!* hard about what has to go into the data file.
2. *Map* the required fields into the NeXus hierarchy
3. *Describe* this map in a NXDL file
4. *Standardize* your definition through communication with the NIAC

3.3.2 Step 1: *Think!* hard about data

This is actually the hard bit. There are two things to consider:

1. What has to go into the data file?
2. What is the normal plot for this type of data?

For the first part, one of the NeXus guiding principles gives us - Guidance! 'A NeXus file must contain all the data necessary for standard data analysis.'

Not more and not less for an application definition. Of course the definition of *standard* data for analysis or a *standard* plot depends on the science and the type of data being described. Consult senior scientists in the field about this is if you are unsure. Perhaps you must call an international meeting with domain experts to haggle that out. When considering this, people tend to put in everything which might come up. This is not the way to go.

A key test question is: Is this data item necessary for common data analysis? Only these necessary data items belong in an application definition.

The purpose of an application definition is that an author of upstream software who consumes the file can expect certain data items to be there at well defined places. On the other hand if there is a development in your field which analyzes data in a novel way and requires more data to do it, then it is better to err towards the side of more data.

Now for the case of WONI, the standard data analysis is either Rietveld refinement or profile analysis. For both purposes, the kind of radiation used to probe the sample (for WONI, neutrons), the wavelength of the radiation, the monitor (which tells us how long we counted) used to normalize the data, the counts and the two theta angle of each detector element are all required. Usually, it is desirable to know what is being analyzed, so some metadata would be nice: a title, the sample name and the sample temperature. The data typically being plotted is two theta against counts, as shown in Figure 3.2 above. Summarizing, the basic information required from WONI is given in Table 3.1.

Table 3.1: basic information required from WONI

`title` of measurement
sample `name`
sample `temperature`
`monitor`
type of radiation `probe`
`wavelength` of radiation incident on sample
`two theta` of detector elements
`counts` for each detector element

If you start to worry that this is too little information, hold on, the section on Using an Application Definition (Section 3.3.7) will reveal the secret how to go from an application definition to a practical file.

3.3.3 Step 2: *Map* Data into the NeXus Hierarchy

This step is actually easier then the first one. We need to map the data items which were collected in Step 1 into the NeXus hierarchy. A NeXus file hierarchy starts with an `NXentry` group. At this stage it is advisable to pull up the base class definition for `NXentry` and study it. The first thing you might notice is that `NXentry` contains a field named `title`. Reading the documentation, you quickly realize that this is a good place to store our title. So the first mapping has been found.

```
title = /NXentry/title
```

Note
In this example, the mapping descriptions just contain the path strings into the NeXus file hierarchy with the class names of the groups to use. As it turns out, this is the syntax used in NXDL link specifications. How convenient!

Another thing to notice in the NXentry base class is the existence of a group of class NXsample. This looks like a great place to store information about the sample. Studying the NXsample base class confirms this view and there are two new mappings:

```
sample name = /NXentry/NXsample/name
sample temperature = /NXentry/NXsample/temperature
```

Scanning the NXentry base class further reveals there can be a NXmonitor group at this level. Looking up the base class for NXmonitor reveals that this is the place to store our monitor information.

```
monitor = /NXentry/NXmonitor/data
```

For the other data items, there seem to be no solutions in NXentry. But each of these data items describe the instrument in more detail. NeXus stores instrument descriptions in the /NXentry/NXinstrument branch of the hierarchy. Thus, we continue by looking at the definition of the NXinstrument base class. In there we find further groups for all possible instrument components. Looking at the schematic of WONI (Figure 3.1), we realize that there is a source, a monochromator and a detector. Suitable groups can be found for these components in NXinstrument and further inspection of the appropriate base classes reveals the following further mappings:

```
probe = /NXentry/NXinstrument/NXsource/probe
wavelength = /NXentry/NXinstrument/NXcrystal/wavelength
two theta of detector elements = /NXentry/NXinstrument/NXdetector/polar angle
counts for each detector element = /NXentry/NXinstrument/NXdetector/data
```

Thus we mapped all our data items into the NeXus hierarchy! What still needs to be done is to decide upon the content of the NXdata group in NXentry. This group describes the data necessary to make a quick plot of the data. For WONI this is counts versus two theta. Thus we add this mapping:

```
two theta of detector elements = /NXentry/NXdata/polar angle
counts for each detector element = /NXentry/NXdata/data
```

The full mapping of WONI data into NeXus is documented in Table 3.2.

Table 3.2: Full mapping of WONI data into NeXus

WONI data	NeXus path
title of measurement	/NXentry/title
sample name	/NXentry/NXsample/name
sample temperature	/NXentry/NXsample/temperature
monitor	/NXentry/NXmonitor/data
type of radiation probe	/NXentry/MXinstrument/NXsource/probe
wavelength of radiation incident on sample	/NXentry/MXinstrument/NXcrystal/wavelength
two theta of detector elements	/NXentry/NXinstrument/NXdetector/polar_angle
counts for each detector element	/NXentry/NXinstrument/NXdetector/data
two theta of detector elements	/NXentry/NXdata/polar_angle
counts for each detector element	/NXentry/NXdata/data

Looking at this one might get concerned that the two theta and counts data is stored in two places and thus duplicated. Stop worrying, this problem is solved at the NeXus API level. Typically NXdata will only hold links to the corresponding data items in /NXentry/NXinstrument/NXdetector.

In this step problems might occur. The first is that the base class definitions contain a bewildering number of parameters. This is on purpose: the base classes serve as dictionaries which define names for everything which possibly can occur. You do not have to give all that information. The key question is, as already said, What is required for typical data analysis for this type of application? You might also be unsure how to correctly store a particular data item. In such a case, contact the NIAC for help. Another problem which can occur is that you require to store information for which there is no name in one of the existing base classes or you have a new instrument component for which there is no base class alltogether. In such a case, please feel free to contact the NIAC with a suggestion for an extension of the base classes in question.

3.3.4 Step 3: *Describe* this map in a NXDL file

This is even easier. Some XML editing is necessary. Fire up your XML editor of choice and open a file. If your XML editor supports XML schema while editing XML, it is worth to load `nxdl.xsd`. Now your XML editor can help you to create a proper NXDL file. As always, the start is an empty template file. This looks like Example 3.1. This is just the basic XML for a NXDL definition. It is advisable to change some of the documentation strings.

Example 3.1 NXDL template file

```
<?xml version="1.0" encoding="UTF-8"?>
<!--
# NeXus - Neutron and X-ray Common Data Format
#
# Copyright (C) 2008-2012 NeXus International Advisory Committee (NIAC)
#
# This library is free software; you can redistribute it and/or
# modify it under the terms of the GNU Lesser General Public
# License as published by the Free Software Foundation; either
# version 3 of the License, or (at your option) any later version.
#
# This library is distributed in the hope that it will be useful,
# but WITHOUT ANY WARRANTY; without even the implied warranty of
# MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See the GNU
# Lesser General Public License for more details.
#
# You should have received a copy of the GNU Lesser General Public
# License along with this library; if not, write to the Free Software
# Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA  02111-1307  USA
#
# For further information, see http://www.nexusformat.org

########### SVN repository information ###################
# $Date: 2012-03-07 00:31:05 +0000 (Wed, 07 Mar 2012) $
# $Author: Pete Jemian $
# $Revision: 1061 $
# $HeadURL: file:///isis/svn/nexus/definitions/branches/docbook/manual/examples/ ↩
    NX__template__.nxdl.xml $
# $Id: NX__template__.nxdl.xml 1061 2012-03-07 00:31:05Z Pete Jemian $
########### SVN repository information ###################
-->
<definition name="NX__template__" extends="NXobject" type="group"
    category="application"
    xmlns="http://definition.nexusformat.org/nxdl/3.1"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xsi:schemaLocation="http://definition.nexusformat.org/nxdl/3.1 ../nxdl.xsd"
    version="1.0b"
    >
    <doc>template for a NXDL application definition</doc>
</definition>
```

For example, copy and rename the file to `NXwoni.nxdl.xml`. Then, locate the XML root element `definition` and change

the `name attribute` (the XML shorthand for this attribute is `/definition/@name`) to `NXwoni`. Change the `doc` as well. Also consider keeping track of `/definition/@version` as suits your development of this NXDL file.

The next thing which needs to be done is adding groups into the definition. A group is defined by some XML like this:

```
<group type="NXdata">

</group>
```

The type is the actual NeXus base class this group belongs to. Optionally a `name` attribute may be given (default is `data`).

Next, one needs to include data items too. The XML for such a data item looks like this:

```
<field name="polar_angle" type="NX_FLOAT units="NX_ANGLE">
  <doc>Link to polar angle in /NXentry/NXinstrument/NXdetector</doc>
  <dimensions rank="1">
    <dim index="1" value="ndet"/>
  </dimensions>
</field>
```

The meaning of the `name` attribute is intuitive, the `type` can be looked up in the relevant base class definition. A `field` definition can optionally contain a `doc` element which contains a description of the data item. The `dimensions` entry specifies the dimensions of the data set. The `size` attribute in the dimensions tag sets the rank of the data, in this example: `rank="1"`. In the `dimensions` group there must be *rank* `dim` fields. Each `dim` tag holds two attributes: `index` determines to which dimension this tag belongs, the `1` means the first dimension. The `value` attribute then describes the size of the dimension. These can be plain integers, variables, such as in the example `ndet` or even expressions like `tof+1`.

Thus a NXDL file can be constructed. The full NXDL file for the WONI example is given in Section 3.3.6. Clever readers may have noticed the strong similarity between our working example `NXwoni` and `NXmonopd` since they are essentially identical. Give yourselves a cookie if you spotted this.

3.3.5 Step 4: *Standardize* with the NIAC

Basically you are done. Your first application definition for NeXus is constructed. In order to make your work a standard for that particular application type, some more steps are required:

- Send your application definition to the NIAC for review
- Correct your definition per the comments of the NIAC
- Cure and use the definition for a year
- After a final review, it becomes the standard

The NIAC must review an application definition before it is accepted as a standard. The one year curation period is in place in order to gain practical experience with the definition and to sort out bugs from Step 1. In this period, data shall be written and analyzed using the new application definition.

3.3.6 Full listing of the WONI Application Definition

```
<?xml version="1.0" encoding="UTF-8"?>
<?xml-stylesheet type="text/xsl" href="nxdlformat.xsl" ?>
<!--
# NeXus - Neutron and X-ray Common Data Format
#
# Copyright (C) 2008-2012 NeXus International Advisory Committee (NIAC)
#
# This library is free software; you can redistribute it and/or
# modify it under the terms of the GNU Lesser General Public
# License as published by the Free Software Foundation; either
```

```
# version 3 of the License, or (at your option) any later version.
#
# This library is distributed in the hope that it will be useful,
# but WITHOUT ANY WARRANTY; without even the implied warranty of
# MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See the GNU
# Lesser General Public License for more details.
#
# You should have received a copy of the GNU Lesser General Public
# License along with this library; if not, write to the Free Software
# Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA  02111-1307  USA
#
# For further information, see http://www.nexusformat.org

########### SVN repository information ###################
# $Date: 2012-03-06 15:00:33 +0000 (Tue, 06 Mar 2012) $
# $Author: Pete Jemian $
# $Revision: 1060 $
# $HeadURL: file:///isis/svn/nexus/definitions/trunk/applications/NXmonopd.nxdl.xml $
# $Id: NXmonopd.nxdl.xml 1060 2012-03-06 15:00:33Z Pete Jemian $
########### SVN repository information ###################
-->
<definition name="NXmonopd" extends="NXobject" type="group"
    category="application"
    xmlns="http://definition.nexusformat.org/nxdl/@NXDL_RELEASE@"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xsi:schemaLocation="http://definition.nexusformat.org/nxdl/@NXDL_RELEASE@ ../nxdl.xsd"
    version="1.0b"
    svnid="$Id: NXmonopd.nxdl.xml 1060 2012-03-06 15:00:33Z Pete Jemian $">
    <doc> Monochromatic Neutron and X-Ray Powder Diffraction. Instrument definition for a powder
        diffractometer at a monochromatic neutron or X-ray beam. This is both suited for a powder
        diffractometer with a single detector or a powder diffractometer with a position sensitive
        detector. </doc>
    <group type="NXentry" name="entry">
        <field name="title"/>
        <field name="start_time" type="NX_DATE_TIME"/>
        <field name="definition">
            <doc> Official NeXus NXDL schema to which this file conforms </doc>
            <enumeration>
                <item value="NXmonopd"/>
            </enumeration>
        </field>
        <group type="NXinstrument">
            <group type="NXsource">
                <field name="type"/>
                <field name="name"/>
                <field name="probe">
                    <enumeration>
                        <item value="neutron"/>
                        <item value="x-ray"/>
                        <item value="electron"/>
                    </enumeration>
                </field>
            </group>
            <group type="NXcrystal">
                <field name="wavelength" type="NX_FLOAT" units="NX_WAVELENGTH">
                    <doc>Optimum diffracted wavelength</doc>
                    <dimensions rank="1">
                        <dim index="1" value="i"/>
                    </dimensions>
```

```
                </field>
            </group>
            <group type="NXdetector">
                <field name="polar_angle" type="NX_FLOAT" axis="1">
                  <doc>where ndet = number of detectors</doc>
                  <dimensions rank="1">
                    <dim index="1" value="ndet" />
                  </dimensions>
                </field>
                <field name="data" type="NX_INT" signal="1">
                  <doc>
                    detector signal (usually counts) are already
                    corrected for detector efficiency
                  </doc>
                  <dimensions rank="1">
                    <dim index="1" value="ndet" />
                  </dimensions>
                </field>
            </group>
        </group>
        <group type="NXsample">
            <field name="name">
                <doc>Descriptive name of sample</doc>
            </field>
            <field name="rotation_angle" type="NX_FLOAT" units="NX_ANGLE">
                <doc> Optional rotation angle for the case when the powder diagram has been ↩
                     obtained
                    through an omega-2theta scan like from a traditional single detector  ↩
                        powder
                    diffractometer </doc>
            </field>
        </group>
        <group type="NXmonitor">
            <field name="mode">
                <doc>Count to a preset value based on either clock time (timer) or received ↩
                     monitor
                    counts (monitor).</doc>
                <enumeration>
                    <item value="monitor"/>
                    <item value="timer"/>
                </enumeration>
            </field>
            <field name="preset" type="NX_FLOAT">
                <doc>preset value for time or monitor</doc>
            </field>
            <field name="integral" type="NX_FLOAT" units="NX_ANY">
                <doc>Total integral monitor counts</doc>
            </field>
        </group>
        <group type="NXdata">
            <link name="polar_angle" target="/NXentry/NXinstrument/NXdetector/polar_angle">
                <doc>Link to polar angle in /NXentry/NXinstrument/NXdetector</doc>
            </link>
            <link name="data" target="/NXentry/NXinstrument/NXdetector/data">
                <doc>Link to data in /NXentry/NXinstrument/NXdetector</doc>
            </link>
        </group>
    </group>
</definition>
```

3.3.7 Using an Application Definition

The application definition is like an interface for your data file. In practice files will contain far more information. For this, the extendable capability of NeXus comes in handy. More data can be added, and upstream software relying on the interface defined by the application definition can still retrieve the necessary information without any changes to their code.

NeXus application definitions only standardize classes. You are free to decide upon names of groups, subject to them matching regular expression for NeXus name attributes (see the regular expression pattern for NXDL group and field names in). Note the length limit of 63 characters imposed by HDF5. Please use sensible, descriptive names and separate multi worded names with underscores.

Something most people wish to add is more metadata, for example in order to index files into a database of some sort. Go ahead, do so, if applicable, scan the NeXus base classes for standardized names. For metadata, consider to use the `NXarchive` definition. In this context, it is worth to mention that a practical NeXus file might adhere to more then one application definition. For example, WONI data files may adhere to both the `NXmonopd` and `NXarchive` definitions. The first for data analysis, the second for indexing into the database.

Often, instrument scientists want to store the complete state of their instrument in data files in order to be able to find out what went wrong if the data is unsatisfactory. Go ahead, do so, please use names from the NeXus base classes.

Site policy might require you to store the names of all your bosses up to the current head of state in data files. Go ahead, add as many `NXuser` classes as required to store that information. Knock yourselves silly over this.

Your Scientific Accounting Department (SAD) may ask of you the preposterous; to store billing information into data files. Go ahead, do so if your judgment allows. Just do not expect the NIAC to provide base classes for this and do not use the prefix NX for your classes.

In most cases, NeXus files will just have one `NXentry` class group. But it may be required to store multiple related data sets of the results of data analysis into the same data file. In this case create more entries. Each entry should be interpretable standalone, i.e. contain all the information of a complete `NXentry` class. Please keep in mind that groups or data items which stay constant across entries can always be linked in.

3.4 Processed Data

Data reduction and analysis programs are encouraged to store their results in NeXus data files. As far as the necessary, the normal NeXus hierarchy is to be implemented. In addition, processed data files must contain a `NXprocess` group. This group, that documents and preserves data provenance, contains the name of the data processing program and the parameters used to run this program in order to achieve the results stored in this entry. Multiple processing steps must have a separate entry each.

Chapter 4

Brief history of the NeXus format

Two things to note about the development and history of NeXus:

- All efforts on NeXus have been voluntary except for one year when we had one full-time worker.
- The NIAC has already discussed many matters related to the format.

June 1994 Mark Koennecke (then ISIS, now PSI) made a proposal using netCDF[1] for the European neutron scattering community while working at ISIS

August 1994 Jonathan Tischler (ORNL) proposed an HDF-based format[2] as a standard for data storage at APS

October 1994 Ray Osborn convened a series of three workshops called *SoftNeSS* (http://www.neutron.anl.gov/softness). In the first meeting, Mark Koennecke and Jon Tischler were invited to meet with representatives from all the major U.S. neutron scattering laboratories at Argonne National Laboratory to discuss future software development for the analysis and visualization of neutron data. One of the main recommendations of *SoftNeSS'94* was that a common data format should be developed.

September 1995 At *SoftNeSS 1995* (at NIST), three individual data format proposals by Przemek Klosowski (NIST), Mark Koennecke (then ISIS), and Jonathan Tischler (ORNL and APS/ANL) were joined to form the basis of the current NeXus format. At this workshop, the name *NeXus* was chosen.

August 1996 The HDF-4 API is quite complex. Thus a NeXus Abstract Programmer Interface (NAPI) was released which simplified reading and writing NeXus files.

October 1996 At *SoftNeSS 1996* (at ANL), after reviewing the different scientific data formats discussed, it was decided to use HDF-4 as it provided the best grouping support. The basic structure of a NeXus file was agreed upon. the various data format proposals were combined into a single document by Przemek Klosowski (NIST), Mark Koennecke (then ISIS), Jonathan Tischler (ORNL and APS/ANL), and Ray Osborn (IPNS/ANL) coauthored the first proposal for the NeXus scientific data standard.[3]

July 1997 SINQ at PSI started writing NeXus files to store raw data.

Summer 2001 MLNSC at LANL started writing NeXus files to store raw data

September 2002 NeXus API version 2.0.0 is released. This version brought support for the new version of HDF, HDF-5, released by the HDF group. HDF-4 imposed limits on file sizes and the number of objects in a file. These issues were resolved with HDF-5. The NeXus API abstracted the difference between the two physical file formats away form the user.

June 2003 Przemek Klosowski, Ray Osborn, and Richard Riedel received the only known grant explicitly for working on NeXus from the Systems Integration for Manufacturing Applications (SIMA) program of the National Institute of Standards and Technology (NIST). The grant funded a person for one year to work on community wide infrastructure in NeXus.

[1] http://wiki.nexusformat.org/images/b/b8/European-Formats.pdf
[2] http://wiki.nexusformat.org/images/d/d5/Proposed_Data_Standard_for_the_APS.pdf
[3] http://wiki.nexusformat.org/images/9/9a/NeXus_Proposal.pdf

October 2003 In 2003, NeXus had arrived at a stage where informal gatherings of a group of people were no longer good enough to oversee the development of NeXus. This lead to the formation of the NeXus International Advisory Committee (NIAC) which strives to include representatives of all major stake holders in NeXus. A first meeting was held at CalTech. Since 2003, the NIAC meets every year to discuss all matters NeXus.

July 2005 The community asked the NeXus team to provide an ASCII based physical file format which allows them to edit their scientific results in emacs. This lead to the development of a XML NeXus physical format. This was released with NeXus API version 3.0.0.

May 2007 NeXus API version 4.0.0 is released with broader support for scripting languages and the feauture to link with external files.

October 2007 NeXus API version 4.1.0 is released with many bug-fixes.

October 2008 NXDL: NeXus Definition Language is defined. Previously, NXDL NeXus used another XML format, meta-DTD, for defining base classes and application definitions. There were several problems with meta-DTD, the biggest one being that it was not easy to validate against it. NXDL was designed to circumvent these problems and all current base classes and application definitions were ported into the NXDL.

April 2009 NeXus API version 4.2.0 is released with additional C++, IDL, and python/numpy interfaces.

September 2009 NXDL and draft `NXsas` presented to canSAS at SAS2009 conference

January 2010 NXDL presented to ESRF HDF5 workshop on hyperspectral data

Chapter 5

NeXus Community

NeXus began as a group of scientists with the goal of defining a common data storage format to exchange experimental results and to exchange ideas about how to analyze them.

The NeXus Scientific Community provides the scientific data, advice, and continued involvement with the NeXus standard. NeXus provides a forum for the scientific community to exchange ideas in data storage through the NeXus wiki.

The NeXus International Advisory Committee supervises the development and maintenance of the NeXus common data format for neutron, x-ray, and muon science. The NIAC supervises a technical committee to oversee the NeXus Application Programmer Interface (NAPI) and the NeXus class definitions.

There are several mechanisms in place in order to coordinate the development of NeXus with the larger community.

5.1 NIAC: The NeXus International Advisory Committee

The purpose of the NeXus International Advisory Committee (NIAC)[1] is to supervise the development and maintenance of the NeXus common data format for neutron, x-ray, and muon science. This purpose includes, but is not limited to, the following activities.

1. To establish policies concerning the definition, use, and promotion of the NeXus format.
2. To ensure that the specification of the NeXus format is sufficiently complete and clear for its use in the exchange and archival of neutron, x-ray, and muon data.
3. To receive and examine all proposed amendments and extensions to the NeXus format. In particular, to ratify proposed instrument and group class definitions, to ensure that the data structures conform to the basic NeXus specification, and to ensure that the definitions of data items are clear and unambiguous and conform to accepted scientific usage.
4. To ensure that documentation of the NeXus format is sufficient, current, and available to potential users both on the internet and in other forms.
5. To coordinate with the developers of the NeXus Application Programming Interface to ensure that it supports the use of the NeXus format in the neutron, x-ray, and muon communities, and to promote other software development that will benefit users of the NeXus format.
6. To coordinate with other organizations that maintain and develop related data formats to ensure maximum compatibility.

The committee will meet at least once every other calendar year according to the following plan:

- In years coinciding with the NOBUGS series of conferences (once every two years), members of the entire NIAC will meet as a satellite meeting to NOBUGS, along with interested members of the community.
- In intervening years, the executive officers of the NIAC will attend, along with interested members of the NIAC. This is intended to be a working meeting with a small group.

[1] For more details about the NIAC constitution, procedures, and meetings, refer to the NIAC wiki page: http://www.nexusformat.org/NIAC The members of the NIAC may be reached by email: NIAC <nexus-committee@nexusformat.org>

5.2 NeXus Mailing Lists

There are several mailing lists associated with NeXus.

NeXus Mailing List We invite anyone who is associated with neutron and/or X-ray synchrotron scattering and who wishes to be involved in the development and testing of the NeXus format to subscribe to this list. It is for the free discussion of all aspects of the design and operation of the NeXus format.

List Address: nexus@nexusformat.org nexus@nexusformat.org
Subscriptions: http://lists.nexusformat.org/mailman/listinfo/nexus
Archive: http://lists.nexusformat.org/pipermail/nexus

NeXus International Advisory Committee (NIAC) Mailing List This list contains discussions of the NeXus International Advisory Committee (NIAC), which oversees the development of the NeXus data format. Its members represent many of the major neutron and synchrotron scattering sources in the world. Membership and posting to this list are confined to the committee members, but the archives are public.

List Address: nexus-committee@nexusformat.org
Subscriptions: http://lists.nexusformat.org/mailman/listinfo/nexus-committee
Archive: http://lists.nexusformat.org/pipermail/nexus-committee

NeXus Developers Mailing List This mailing list is for discussions concerning the technical development of NeXus (the Definitions, NXDL, and the NeXus Application Program Interface).

List Address: nexus-developers@nexusformat.org
Subscriptions: http://lists.nexusformat.org/mailman/listinfo/nexus-developers
Archive: http://lists.nexusformat.org/pipermail/nexus-developers

Subversion (http://subversion.tigris.org) is the revision control system used by the NeXus developers.

TRAC (http://trac.edgewall.org) is the issue tracking and bug reporting system used by the NeXus developers.

NeXus Code Subversion Mailing List Members of this list will receive an email whenever a commit is made to the NeXus code repository. This list cannot be posted to - all questions should instead be sent to the NeXus Developers Mailing List (nexus-developers@nexusformat.org).

List Address: nexus-code-svn@nexusformat.org
Subscriptions: http://lists.nexusformat.org/mailman/listinfo/nexus-code-svn
Archive: http://lists.nexusformat.org/pipermail/nexus-code-svn

NeXus Code Tickets Mailing List Members of this list will receive an email whenever a ticket (bug/issue/task) associated with NeXus code library development is modified on the Nexus *code* TRAC server. The list of ticket updates and subversion changesets is available on the *code* repository TRAC timeline. This list cannot be posted to - see the section on Issue Reporting.

List Address: nexus-code-tickets@nexusformat.org
Subscriptions: http://lists.nexusformat.org/mailman/listinfo/nexus-code-tickets
Archive: http://lists.nexusformat.org/pipermail/nexus-code-tickets
TRAC Timeline: http://trac.nexusformat.org/code/report/1

NeXus Definitions Subversion Mailing List Members of this list will receive an email whenever a commit is made to the NeXus definitions repository. This list cannot be posted to - all questions should instead be sent to the NeXus Developers Mailing List (nexus-developers@nexusformat.org).

List Address: nexus-definitions-svn@nexusformat.org
Subscriptions: http://lists.nexusformat.org/mailman/listinfo/nexus-definitions-svn
Archive: http://lists.nexusformat.org/pipermail/nexus-definitions-svn

NeXus Definitions Tickets Mailing List Members of this list will receive an email whenever a ticket (bug/issue/task) associated with NeXus definitions development is modified on the Nexus *definitions* TRAC server. The list of ticket updates and subversion changesets is available on the *definitions* repository TRAC timeline. This list cannot be posted to - see the section on Issue Reporting.

List Address: nexus-definitions-tickets@nexusformat.org
Subscriptions: http://lists.nexusformat.org/mailman/listinfo/nexus-definitions-tickets
Archive: http://lists.nexusformat.org/pipermail/nexus-definitions-tickets
TRAC Timeline: http://trac.nexusformat.org/definitions/report/1

5.3 NeXus Subversion Repositories

NeXus NXDL class definitions (both base classes and instruments) and the NeXus code library source are held in a subversion repository. The repository is world readable and though you can browse the NeXus code library and applications or NeXus NXDL class definitions repositories directly, a better looking interface is provided by the ViewVC or TRAC browsers.

Browse the NeXus code (library and applications) repository using ViewVC or TRAC
Browse NeXus definitions (NXDL classes) repository using ViewVC or TRAC

The repository can alse be interrogated for recent updates via a query form

```
http://svn.nexusformat.org/viewvc/NeXusCode/trunk/?view=queryform
```

For example, show me all changes in the last month for the code (library and applications) repository

```
http://svn.nexusformat.org/viewvc/NeXusCode/trunk/?view=query&date=month&limit_changes=100
```

or Definition repository

```
http://trac.nexusformat.org/definitions/timeline?daysback=30
```

If you wish to receive an email when a change is made to the repository you should join the appropriate Mailing Lists.

Table 5.1: TRAC RSS feed

Alternatively, you can use an RSS feed to keep abreast of changes. TRAC provides a link to its RSS feed on pages with an orange *XML RSS Feed* icon at their foot such as:	

There are pages that show the subversion repository activity in a timeline format or a tabular (revision log) format.

code (library and applications) repository timeline `http://trac.nexusformat.org/code/timeline`

definitions repository timeline `http://trac.nexusformat.org/definitions/timeline`

code repository revision log `http://trac.nexusformat.org/code/log`

definitions repository revision log `http://trac.nexusformat.org/definitions/log`

5.3.1 Login

To update files in these repositories you will need to use a subversion client such as `TortoiseSVN/`[2] for Microsoft Windows or `svn` for command-line shells and also provide your NeXus Wiki username and password. Note that for subversion write access:

- If your Wiki username contains a space, write it with a space (i.e. do not replace the space with an _ as is done in WIKI URLs)
- You cannot use a *temporary password* (i.e. one that was emailed to you in response to a request). You must first log into MediaWiki with the temporary password and then go to account NeXus wiki Preferences and change the password.
- Your Wiki account must have an email address associated with it and this address must have been validated. To provide and/or validate your email address, log in and go to your account NeXus wiki Preferences. section.
- If you have login problems and have not changed your WIKI password since 20th October 2006, please go to the NeXus wiki login page and request to be emailed a new password. To synchronise TRAC/Subversion/MediaWiki required some changes to the authentication system which will have invalidated passwords set prior to that date.

[2] `http://tortoisesvn.tigris.org/`

Here are the URLs to access the subversion repositories as a developer:

code for library/applications

```
https://svn.nexusformat.org/code/trunk
```

definitions for NXDL classes

```
https://svn.nexusformat.org/definitions/trunk
```

checkout the code trunk

```
svn co --username "My WIKI Username" https://svn.nexusformat.org/code/trunk nexus_code
```

Please report any problems via the Issue Reporting system.

5.3.2 Committing Changes

As well as needing a valid account, you will not be able to check-in changes unless you indicate (in the log message attached to the commit) which current issues on the Issue Reporting system the changes either fix or refer to. This is done by enclosing special phrases in the commit message of the form:

```
command #1
command #1, #2
command #1 & #2
command #1 and #2
```

where `command` is one of the commands detailed below and `#1` means *issue number 1* on the system, etc. You can have more then one command in a message. The following commands are supported and there is more then one spelling for each command (to make this as user-friendly as possible):

`closes, fixes` The specified issue numbers are closed with the contents of this commit message being added to it.

`references, refs, addresses, re` The specified issue numbers are left in their current status, but the contents of this commit message are added to their notes.

For example, the commit message

```
Changed blah and foo to do this or that. Fixes #10 and #12, and refs #12.
```

This will close issues #10 and #12, and add a note to #12 on the Issue Reporting system. For a list of current issues, see:

Active tickets for the NeXus code library: http://trac.nexusformat.org/code/report/1
Active tickets for NeXus definitions: http://trac.nexusformat.org/definitions/report/1

5.3.3 URLs described in this section

Many Uniform Resource Locators (URLs) have been used in this section. This is a table describing them.

Subversion revision management software

http://subversion.tigris.org/

ViewVC versions control repository viewing software

http://www.viewvc.org/

TRAC issue management software

http://trac.edgewall.org

TortoiseSVN, Windows subversion client

http://tortoisesvn.tigris.org/

NeXus code (library and applications) subversion repository

http://svn.nexusformat.org/code/

NeXus definitions subversion repository

http://svn.nexusformat.org/definitions/

ViewVC view of NeXus code (library and applications) repository

http://svn.nexusformat.org/viewvc/NeXusCode

ViewVC view of NeXus definitions repository

http://svn.nexusformat.org/viewvc/NeXusDefinitions

TRAC view of NeXus code (library and applications) repository

http://trac.nexusformat.org/code/browser

NeXus code (library and applications) revision log

http://trac.nexusformat.org/code/log

Active tickets for the NeXus code repository

http://trac.nexusformat.org/code/report/1

NeXus code repository timeline

http://trac.nexusformat.org/code/timeline

TRAC view of NeXus definitions repository

http://trac.nexusformat.org/definitions/browser

NeXus definitions revision log

http://trac.nexusformat.org/definitions/log

Active tickets for NeXus definitions

http://trac.nexusformat.org/definitions/report/1

NeXus definitions repository timeline

http://trac.nexusformat.org/definitions/timeline

NeXus code repository (password required)

https://svn.nexusformat.org/code/trunk

NeXus definitions repository (password required)

https://svn.nexusformat.org/definitions/trunk

5.4 NeXus Issue Reporting

NeXus is using TRAC[3] for problem/issue reporting. The issue reports (see *View current issues* below) are used to guide the NeXus developers in resolving problems as well as implementing new features. As such, the TRAC tickets for the *code* and *definitions* repositories form the basis of a *roadmap* for NeXus. You can browse issues without logging on, but to report issues you will need to login using your NeXus WIKI username and password (the subversion login notes mentioned for write access to the Subversion Server apply to TRAC login, too).

Whenever an update is made to a ticket, a message is also posted to the appropriate ticket mailing list.

[3] http://trac.edgewall.org

5.4.1 NeXus Code (Library and Applications)

Report a new issue: http://trac.nexusformat.org/code

View current issues: http://trac.nexusformat.org/code/report/1

Archive of ticket update emails: http://lists.nexusformat.org/pipermail/nexus-code-tickets/

repository timeline (recent ticket and code changes): http://trac.nexusformat.org/code/timeline

repository roadmap: http://trac.nexusformat.org/code/roadmap

5.4.2 NeXus Definitions (NXDL base classes and application definitions)

Report a new issue: http://trac.nexusformat.org/definitions

View current issues: http://trac.nexusformat.org/definitions/report/1

Archive of ticket update emails: http://lists.nexusformat.org/pipermail/nexus-definitions-tickets/

repository timeline (recent ticket and definition changes): http://trac.nexusformat.org/definitions/timeline

repository roadmap: http://trac.nexusformat.org/definitions/roadmap

Appendix A

Installation

This section descibes how to install the NeXus API and details the requirements. The NeXus API is distributed under the terms of the GNU Lesser Public License.

The source code and binary versions for some popular platforms can be found on http://download.nexusformat.org/kits/. Up to date instructions can be found on the Wiki In case you need help feel free to contact the nexus mailing list.

A.1 Precompiled Binary Installation

A.1.1 Prerequisites

A.1.1.1 HDF5/HDF4

Since NeXus uses HDF as the main underlying binary format, it is necessary first to install the HDF subroutine libraries and include files before compiling the NeXus API. It is not usually necessary to download the HDF source code since precompiled object libraries exist for a variety of operating systems including Windows, Mac OS X, Linux, and various other flavors of Unix. Check the HDF web pages for more information: http://www.hdfgroup.org/

Packages for HDF4 and HDF5 are available for both Fedora (hdf, hdf5, hdf-devel, hdf5-devel) and Ubuntu/Debian (libhdf4g, libhdf5).

A.1.1.2 XML

The NeXus API also supports using XML as the underlying on-disk format. This uses the Mini-XML library, developed by Michael Sweet, which is also available as a precompiled binary library for several operating systems. Check the Mini-XML web pages for more information: http://www.minixml.org/

Packages for MXML are available for both Fedora (mxml, mxml-devel) and Ubuntu/Debian (libmxml1).

A.1.2 Linux RPM Distribution Kits

An installation kit (source or binary) can be downloaded from: http://download.nexusformat.org/kits/

A NeXus binary RPM (nexus-*.i386.rpm) contains ready compiled NeXus libraries whereas a source RPM (nexus-*.src.rpm) needs to be compiled into a binary RPM before it can be installed. In general, a binary RPM is installed using the command

```
rpm -Uvh file.i386.rpm
```

or, to change installation location from the default (e.g. /usr/local) area, using

```
rpm -Uvh --prefix /alternative/directory file.i386.rpm
```

If the binary RPMS are not the correct architecture for you (e.g. you need x86_64 rather than i386) or the binary RPM requires libraries (e.g. HDF4) that you do not have, you can instead rebuild a source RPM (.src.rpm) to generate the correct binary RPM for you machine. Download the source RPM file and then run

```
rpmbuild --rebuild file.src.rpm
```

This should generate a binary RPM file which you can install as above. Be careful if you think about specifying an alternative buildroot for rpmbuild by using `--buildroot` option as the "buildroot" directory tree will get remove (so `--buildroot /` is a really bad idea). Only change buildroot it if the default area turns out not to be big enough to compile the package.

If you are using Fedora, then you can install all the dependencies by typing

```
yum install hdf hdf-devel hdf5 hdf5-devel mxml mxml-devel
```

A.1.3 Microsoft Windows Installation Kit

A Windows MSI based installation kit is available and can be downloaded from:http://download.nexusformat.org/kits/windows/

A.1.4 Mac OS X Installation Kit

An installation disk image (.dmg) can be downloaded from:http://download.nexusformat.org/kits/macosx/

A.2 Source Installation

A.2.1 NeXus Source Code Distribution

The build uses `autoconf` (so autools are required) to determine what features will be available by your system. You must have the *development* libraries installed for all the file backends you want support for (see above). If you intend to build more than the C language bindings, you need to have the respective build support in a place where autoconf will pick them up (i.e. python development files, a Java Development Kit, etc.).

For more information see the README in the toplevel of the source distribution. In case you need help, feel free to contact the developers using the `nexus-developers mailing list`.

Download the appropriate gzipped tar file, unpack it, and run the standard configure procedure from the resulting nexus directory. For example, for version 4.2.1;

```
$ tar zxvf nexus-4.2.1.tar.gz
$ cd nexus-4.2.1
$ ./configure
```

To find out how to customize the installation, e.g., to choose different installation directories, type

```
$ ./configure --help
```

Carefully check the final output of the `configure` run. Make sure all features requested are actually enabled.

```
$ make
$ make install
```

See the README file for further instructions.

A.2.2 Cygwin Kits

HDF4 is not supported under CYGWIN - both HDF5 and MXML are supported and can be downloaded and built as usual. When configuring HDF5 you should explicitly pass a prefix to the configure script to make sure the libraries are installed in a "usual" location i.e.

```
./configure --prefix=/usr/local/hdf5
```

Otherwise you will have to use the `--with-hdf5=/path/to/hdf5` option later when configuring NeXus to tell it where to look for hdf5. After building hdf5, configure and build NeXus using the instructions for source code distribution above.

Appendix B

Verification and validation of files

The intent of verification and validation of files is to ensure, in an unbiased way, that a given file conforms to the relevant specifications. NeXus uses various automated tools to validate files. These tools include conversion of content from HDF to XML and transformation (via XSLT) from XML format to another such as NXDL, XSD, and Schematron. This chapter will first provide an overview of the process, then define the terms used in validation, then describe how multiple base classes or application definitions might apply to a given NeXus data file, and then describe the various validation techniques in more detail. Validation does not check that the data content of the file is sensible; this requires scientific interpretation based on the technique.

Validation is useful to anyone who manipulates or modifies the contents of NeXus files. This includes scientists/users, instrument staff, software developers, and those who might mine the files for metadata. First, the scientist or user of the data must be certain that the information in a file can be located reliably. The instrument staff or software developer must be confident the information they have written to the file has been located and formatted properly. At some time, the content of the NeXus file may contribute to a larger body of work such as a metadata catalog for a scientific instrument, a laboratory, or even an entire user facility.

B.1 Overview

NeXus files adhere to a set of rules and can be tested against these rules for compliance. The rules are implemented using standard tools and can themselves be tested to verify compliance with the standards for such definitions. Validation includes the testing of both NeXus data files and the NXDL specifications that describe the rules.

The rules for writing NeXus data files are different than the rules for writing NeXus class definitions. To validate a NeXus data file, these two rule sets must eventually merge, as shown in the next figure. The data file (either HDF4, HDF5, or XML) is first converted into an internal format to facilitate validation, including data types, array dimensions, naming, and other items. Most of the data is not converted since data validation is non-trivial. Also note that the units are not validated. All the NXDL files are converted into a single Schematron file (again, internal use for validation) only when NXDL revisions are checked in to the NeXus definitions repository as NXDL changes are not so frequent.

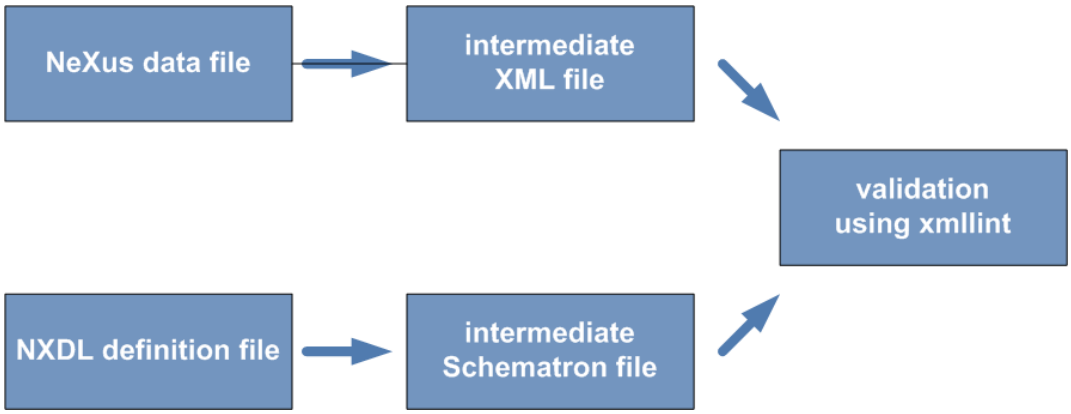

Figure B.1: Flowchart of the NeXus validation process.

NeXus data files

NeXus data files (also known as NeXus data file instances) are validated to ensure the various parts of the data file are arranged according to the governing NXDL specifications used in that file instance.

Note
Since NeXus has several rules that are quite difficult to apply in either XSD or Schematron, direct validation of data files using standard tools is not possible. To validate NeXus data files, it is necessary to use `nxvalidate`.

NeXus Definition Language (NXDL) specification files

NXDL files are validated to ensure they adhere to the rules for writing NeXus base classes and application definitions.

B.2 Definitions of these terms

Let's be clear about some terms used in this section.

HDF

Hierarchical Data Format from The HDF Group. NeXus data files using HDF may be stored in either version 4 (HDF4) or version 5 (HDF5). New NeXus HDF files should only use HDF5. The preferred file extensions (but not required) include `.hdf`, `.h5`, `.nxs`, and `.nx5`.

NXDL

NeXus Definition Language files define the spcifications for NeXus base classes, application definitions, and contributed classes and definitions. It is fully described in the NXDL chapter in Volume II of this documentation.

Schematron

Schematron.[1] is an alternative to XSD and is used to validate the content and structure of an XML file. NeXus uses Schematron internally to validate data files.

Validation

File validation is the comparison of file contents, in an unbiased way, with the set of rules that define the structure of such files.

[1] http://www.schematron.com

XML

The eXtensible Markup Language (XML) eXtensible Markup Language (XML).[2] is a standard business tool for the exchange of information. It is broadly supported by a large software library in many languages. NeXus uses XML for several purposes: data files, NXDL definitions, rules, and XSLT transformations.

XSD

XML files are often defined by a set of rules (or *schema*). A common language used to implement these rules is XML Schema (XSD).[3] Fundamentally, XML, XSD, XSLT, and Schematron files are all in XML.

XSLT

XML files can be flexible enough to convert from one set of rules to another. An example is when one company wishes to exchange catalog or production information with another. The XML StyLsheet Transformation (XSLT)[4] (XSLT) language is often used to describe each direction of the conversion of the XML files between the two rule sets.

B.3 NeXus data files may use multiple base classes or application definitions

NeXus data files may have more than one data set or may have multiple instances of just about any base class or even application definitions. The NeXus data file validation is prepared to handle this without any special effort by the provider of the data file.

B.4 Validation techniques

File validation is the process to determine if a given file is prepared consistent with a set of guidelines or rules. In NeXus, there are several different types of files. First, of course, is the data file yet it can be provided in one of several forms: HDF4, HDF5, or XML. Specifications for data files are provided by one or (usually) more NeXus definition files (NXDL, for short). These NXDL files are written in XML and validated by the NXDL specification which is written in the XML Schema (or XSD) language. Thus, automated file verification is available for data files, definition files, and the rules for definition files.

B.4.1 Validation of NeXus data files

Each NeXus data file can be validated against the NXDL rules. (The full suite of NXDL specifications is converted into Schematron rules by an XSLT transformation and then combined into a single file. It is not allowed to have a NeXus base class and also an application definition with the same name since one will override the other in the master Schematron file) The validation is done using Schematron and the `NXvalidate` program. Schematron was selected, rather than XML Schema (XSD) , to permit established rules for NeXus files, especially the rule allowing the nodes within `NXentry` to appear in any order.

The validation process is mainly checking file structure (presence or absence of groups/fields) - it is usually impossible to check the actual data itself, other than confirm that it is of the correct data type (string, float etc.). The only exception is when the NXDL specification is either a fixed value or an enumeration - in which case the data is checked.

During validation, the NeXus data file instance (either HDF or XML) is first converted into an XML file in a form that facilitates validation (e.g with large numeric data removed). Then the XML file is validated by Schematron against the `schema/all.sch` file.

B.4.2 Validation of NeXus Definition Language (NXDL) specification files

Each NXDL file must be validated against the rules that define how NXDL files are to be arranged. The NXDL rules are specified in the form of XML Schema (XSD).

Standard tools (validating editor or command line or support library) can be used to validate any NXDL file. Here's an example using `xmllint` from a directory that contains `nxdl.xsd`, `nxdlTypes.xsd`, and `applications/NXsas.nxdl.xml`:

[2] http://www.w3schools.com/xml
[3] http://www.w3schools.com/xsd
[4] http://www.w3schools.com/xsl

Example B.1 Use of `xmllint` to validate a NXDL specification.

```
xmllint noout --schema nxdl.xsd applications/NXsas.nxdl.xml
```

B.4.3 Validation of the NXDL rules

NXDL rules are specified using the rules of XML Schema (XSD). The XSD syntax of the rules is validated using standard XML file validation tools: either a validating editor (such as *oXygen*, *xmlSpy*, or *eclipse*) or common UNIX/Linux command line tools

Example B.2 Use of `xmllint` to validate the NXDL rules.

```
xmllint --valid  nxdl.xsd
```

The validating editor method is used by the developers while the `xmllint` command line tool is the automated method used by the NeXus definitions subversion repository.

B.4.4 Validation of XSLT files

XSLT transformations are validated using standard tools such as a validating editor or xmllint.

B.4.5 Transformation of NXDL files to Schematron

Schematron[5] is a rule-based language that allows very specific validation of an XML document. Its advantages over using XSD schema are that:

- more specific pattern-based rules based on data content can be written
- full XSLT/XPath expression syntax available for writing validation tests
- error messages can be customised and thus more meaningful
- It is easier to validate documents when entities can occur in any order.

XSD does provide a mechanism for defining a class structure and inheritance, so its usage within NeXus in addition to schematron has not been ruled out. But for a basic validation of file content, schematron looks best.

The NXDL definition files are converted into a set of Schematron rules using the `xslt/nxdl2sch.xsl` XSLT stylesheet. The NeXus instance file (either in XML, HDF4, or HDF5) is turned into a reduced XML validation file. This file is very similar to a pure NeXus XML file, but with additional metadata for dimenions and also with most of the actual numeric data removed.

The validation process then compares the set of Schematron rules against the *reduced XML* validation file. Schematron itself is implemented as a set of XSLT transforms. NeXus includes the Schematron files, as well as the Java based XSLT engine `saxon`.

The java based `nxvalidate` GUI can be run to validate files.

Currently, the structure of the file is validated (i.e. valid names are used at the correct points), but this will be extended to array dimensions and link targets. Error messages are printed about missing mandatory fields, and informational messages are printed about fields that are neither optional or mandatory (in case they are a typing error). Even non-standard names must comply with a set of rules (e.g. no spaces are allowed in names). Enumerations are checked that they conform to an allowed value. The data type is checked and the units will also be checked.

[5] http://www.schematron.com

Appendix C

NeXus Utilities

There are many utilities available to read, browse, write, and use NeXus data files. Some are provided by the NeXus technical group while others are provided by the community. Still, other tools listed here can read or write one of the low-level file formats used by NeXus (HDF4, HDF5, or XML).

C.1 Utilities supplied with NeXus

Most of these utility programs are run from the command line. It will be noted if a program provides a graphical user interface (GUI). Short descriptions are provided here with links to further information, as available.

nxbrowse NeXus Browser

nxconvert Utility to convert a NeXus file into HDF4/HDF5/XML/...

nxdir `NXdir` is a utility for querying a NeXus file about its contents. Full documentation can be found by running this command:

```
nxdir -h
```

nxingest `nxingest` extracts the metadata from a NeXus file to create an XML file according to a mapping file.

The mapping file defines the structure (names and hierarchy) and content (from either the NeXus file, the mapping file or the current time) of the output file. See the man page for a description of the mapping file.

This tool uses the NAPI. Thus, any of the supported formats (HDF4, HDF5 and XML) can be read.

nxsummary Use `nxsummary` to generate summary of a NeXus file.

This program relies heavily on a configuration file. Each `item` tag in the file describes a node to print from the NeXus file. The `path` attribute describes where in the NeXus file to get information from. The `label` attributes is what will be printed when showing the value of the specified field. The optional `operation` attribute provides for certain operations to be performed on the data before printing out the result.

See the source code documentation for more details.

nxtranslate `nxtranslate` is an anything to NeXus converter. This is accomplished by using translation files and a plugin style of architecture where `nxtranslate` can read from new formats as plugins become available. The documentation for `nxtranslate` describes its usage by three types of individuals:

- the person using existing translation files to create NeXus f
- the person creating translation files
- the person writing new `Retrievers`

All of these concepts are discussed in detail in the documentation provided with the source code.

nxvalidate From the source code documentation: 'Utility to convert a NeXus file into HDF4/HDF5/XML/...' Note: this command-line tool is different than the newer Java GUI program: `NXvalidate`.

nxvalidate Java program (in development in 2010) to check any NeXus data file for conformance with the NeXus NXDL-based standard. Note: This Java GUI is different than the command-line tool: `nxvalidate`.

NXdump `NXdump` is temporary wrapper script for `.libs/NXdump` in the NeXus code repository. From the source code documentation: ‘ The NXdump program cannot be directly executed until all the libtool libraries that it depends on are installed. This wrapper script should never be moved out of the build directory. If it is, it will not operate correctly. ’

NXplot An extendable utility for plotting any NeXus file. `NXplot` is an Eclipse-based GUI project in Java to plot data in NeXus files. (The project was started at the first NeXus Code Camp in 2009.)

C.2 Data Analysis

The list of applications below are some of the utilities that have been developed (or modified) to read/write NeXus files as a data format. It is not intended to be a complete list of all available packages.

DAVE (www.ncnr.nist.gov/dave) DAVE is an integrated environment for the reduction, visualization and analysis of inelastic neutron scattering data. It is built using IDL (Interactive Data Language) from ITT Visual Information Solutions.

GDA (www.opengda.org) The GDA project is an open-source framework for creating customised data acquisition and analysis software for science facilities such as neutron and x-ray sources.

Gumtree (http://docs.codehaus.org/display/GUMTREE) Gumtree is an open source project, providing a graphical user interface for instrument status and control, data acquisition and data reduction.

ISAW (ftp.sns.gov/ISAW) The Integrated Spectral Analysis Workbench software project (ISAW) is a Platform-Independent system Data Reduction/Visualization. ISAW can be used to read, manipulate, view, and save neutron scattering data. It reads data from IPNS run files or NeXus files and can merge and sort data from separate measurements.

LAMP (www.ill.eu/data_treat/lamp) LAMP (Large Array Manipulation Program) is designed for the treatment of data obtained from neutron scattering experiments at the Institut Laue-Langevin. However, LAMP is now a more general purpose application which can be seen as a GUI-laboratory for data analysis based on the IDL language.

Mantid (www.mantidproject.org) The Mantid project provides a platform that supports high-performance computing on neutron and muon data. It is being developed as a collaboration between Rutherford Appleton Laboratory and Oak Ridge National Laboratory.

NeXpy (trac.mcs.anl.gov/projects/nexpy) The goal of NeXpy is to provide a simple graphical environment, coupled with Python scripting capabilities, for the analysis of X-Ray and neutron scattering data. (It was decided at the NIAC 2010 meeting that a large portion of this code would be adopted in the future by NeXus and be part of the distribution)

OpenGENIE (www.opengenie.org) A general purpose data analysis and visualisation package primarily developed at the ISIS Facility, Rutherford Appleton Laboratory.

PyMCA (pymca.sourceforge.net) PyMca is a ready-to-use, and in many aspects state-of-the-art, set of applications implementing most of the needs of X-ray fluorescence data analysis. It also provides a toolkit of Python tools for visualization and analysis of energy-dispersive X-ray fluorescence data.

C.3 HDF Tools

Here are some of the generic tools that are available to work with HDF files. In addition to the software listed here there are also APIs for many programming languages that will allow low level programmatic access to the data structures.

Command Line Tools (www.hdfgroup.org/products/hdf5_tools) There are various command line tools that are available from the HDF Group, these are usually shipped with the HDF5 kits but are also available for download separately.

HDFexplorer (www.space-research.org) A data visualization program that reads Hierarchical Data Format files (HDF, HDF-EOS and HDF5) and also netCDF data files.

HDFview (www.hdfgroup.org) A Java based GUI for browsing (and some basic plotting) of HDF files.

IDL (www.ittvis.com) IDL is a high-level technical computing language and interactive environment for algorithm development, data visualization, data analysis, and numeric computation.

IgorPro (www.wavemetrics.com) IGOR Pro is an extraordinarily powerful and extensible scientific graphing, data analysis, image processing and programming software tool for scientists and engineers.

MATLAB (www.mathworks.com) MATLAB is a high-level technical computing language and interactive environment for algorithm development, data visualization, data analysis, and numeric computation.

Appendix D

Frequently Asked Questions

This is a list of commonly asked questions concerning the NeXus data format.

1. *How many facilities use NeXus?*

   This is not easy to say, not all facilities using NeXus actively participate in the committee. Some facilities have reported their adoption status on the `Facilities Wiki page` . Please have a look at this list. Keep in mind that it is not complete.

2. *NeXus files are binary? This is crazy! How am I supposed to see my data?*

   NeXus files are not per se binary. If you use the XML backend the data are stored in a relatively human readable form (see this example). This backend however is only recommended for very small data sets. With the multidimensional data that is routinely recorded on many modern instruments it is very difficult anyway to retrieve useful information on a VT100 terminal. If you want to try, for example `nxbrowse` is a utility provided by the NeXus community that can be very helpful to those who want to inspect their files and avoid graphical applications. For larger data volumes the binary backends used with the appropriate tools are by far superior in terms of efficiency and speed and most users happily accept that after having worked with supersized "human readable" files for a while.

3. *What on-disk file format should I choose for my data?*

   HDF5 is the default file container to use for NeXus data. It is the recommended format for all applications. HDF4 is still supported as a on disk format for NeXus but for new installations preference should be given to HDF5. The XML backend is available for special use cases. Choose this option with care considering the space and speed implications.

4. *Why are the NeXus classes so complicated? I'll never store all that information*

   The NeXus classes are essentially glossaries of terms. If you need to store a piece of information, consult the class definitions to see if it has been defined. If so, use it. It is not compulsory to include every item that has been defined in the base class if it is not relevant to your experiment. On the other hand, a NeXus application definition lists a smaller set of compulsory items that should allow other researchers or software to analyze your data. You should really follow the application definition that corresponds to your experiment to take full advantage of NeXus.

5. *I don't like NeXus. I seems much faster and simpler to develop my own file format. Why should I even consider NeXus?*

   If you consider using an efficient on disk storage format, HDF5 is a better choice than most others. It is fast and efficient and well supported in all main stream programming languages and a fair share of popular analysis packages. The format is so widely used and backed by a big organisation that it will continue to be supported for the foreseeable future. So if you are going to use HDF5 anyway, why not use the NeXus definition to lay out the data in a standardised way? The NeXus community spent years trying to get the standard right and while you will not agree with every single choice they made in the past, you should be able to store the data you have in a quite reasonable way. If you do not comply with NeXus chances are most people will perceive your format as different but not necessarily better than NeXus by any large measure. So it may not be worth the effort. Seriously. If you encounter any problems because the classes are not sufficient to describe your configuration, please contact the NIAC Executive Secretary explaining the problem, and post a suggestion at the relevant class wiki page. Or raise the problem in one of the mailing lists. The NIAC is always willing to consider new proposals.

6. *I want to produce an application definition. How do I go about it?*

   Read the NXDL Tutorial in Creating a NXDL Specification. The procedures for acceptance are defined in the NIAC constitution. [1]

7. *What is the purpose of `NXdata`?*

   `NXdata` contains links to the data stored elsewhere in the `NXentry`. It identifies the default plottable data. This is one of the basic motivations (see Simple plotting) for the NeXus standard. The choice of the name `NXdata` is historic and does not really reflect its function.

8. *How do I identify the plottable data?*

   See the section: Section 2.4.8.

9. *How can I specify reasonable axes for my data?*

   See the section: Section 2.4.

10. *Why aren't `NXsample` and `NXmonitor` groups stored in the `NXinstrument` group?*

    A NeXus file can contain a number of `NXentry` groups, which may represent different scans in an experiment, or sample and calibration runs, etc. In many cases, though by no means all, the instrument has the same configuration so that it would be possible to save space by storing the `NXinstrument` group once and using multiple links in the remaining `NXentry` groups. It is assumed that the sample and monitor information would be more likely to change from run to run, and so should be stored at the top level.

11. *Specifications are complicated and often provide too much information for what I need. Where can I find some good example data files?*

    There are a few checked into the definitions repository. At the moment the selection is quite limited and not very representative. This repository will be edited as more example files become available.

12. *Can I use a NXDL specification to parse a NeXus data file?*

    This should be possible as there is nothing in the NeXus specifications to prevent this but it is not implemented in NAPI. You would need to implement it for yourself. You would be wise to consult the algorithms in the Java version of `NXvalidate` (see NXvalidate (java)) for more details.

13. *Why do I need to specify the `NAPItype`? My programming language does not need that information and I don't care about C and colleagues. Can I leave it out?*

    `NAPItype` is necessary. When implementing the NeXus-XML API we strived to make this as general as HDF and reasonably efficient for medium sized datasets. This is why we store arrays as a large bunch of numbers in C-storage order. And we need the `NAPItype` to figure out the dimensions of the dataset.

14. *Do I have to use the `NAPI` subroutines? Can't I read (or write) the NeXus data files with my own routines?*

    You are not required to use the NAPI to write valid NeXus data files. It is possible to avoid the NAPI to write and read valid NeXus data files. But, the programmer who chooses this path must have more understanding of how the NeXus HDF or XML data file is written. Validation of data files written without the NAPI is strongly encouraged.

15. *I'm using links to place data in two places. Which one should be the data and which one is the link?*

    NeXus uses HDF5 hard links. Both places have pointers to the actual data. That is the way hard links work in HDF5. There is no need for a preference to either location. NeXus defines a `target` attribute to label one directory entry as the source of the data (in this, the link *target*). This has value in only a few situations such as when converting the data from one format to another. By identifying the original in place, duplicate copies of the data are not converted. In HDF, a hard link points to a data object. A soft link points to a directory entry. Since NeXus uses hard links, there is no need to distinguish between two (or more) directory entries that point to the same data.

16. *If I write my data according to the current specification for NXsas (substitute any other application definition), will other software be able to read my data?*

    Yes. NXsas, like other Application Definitions, defines and names the *minimum information* required for analysis or data processing. As long as all the information required by the specification is present, analysis software should be able to process the data. If other information is also present, there is no guarantee that small-angle scattering analysis software will notice.

[1] Refer to the most recent version of the NIAC constitution on the NIAC wiki: `http://www.nexusformat.org/NIAC`

Appendix E

Licenses

The full texts of the software licenses governing this document (FDL) and the example in this document (LGPL) are provided in this appendix.

E.1 FDL: GNU Free Documentation License


```
                GNU Free Documentation License
                 Version 1.3, 3 November 2008


 Copyright (C) 2000, 2001, 2002, 2007, 2008 Free Software Foundation, Inc.
     <http://fsf.org/>
 Everyone is permitted to copy and distribute verbatim copies
 of this license document, but changing it is not allowed.

0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other
functional and useful document "free" in the sense of freedom: to
assure everyone the effective freedom to copy and redistribute it,
with or without modifying it, either commercially or noncommercially.
Secondarily, this License preserves for the author and publisher a way
to get credit for their work, while not being considered responsible
for modifications made by others.

This License is a kind of "copyleft", which means that derivative
works of the document must themselves be free in the same sense.  It
complements the GNU General Public License, which is a copyleft
license designed for free software.

We have designed this License in order to use it for manuals for free
software, because free software needs free documentation: a free
program should come with manuals providing the same freedoms that the
software does.  But this License is not limited to software manuals;
it can be used for any textual work, regardless of subject matter or
whether it is published as a printed book.  We recommend this License
principally for works whose purpose is instruction or reference.


1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that
contains a notice placed by the copyright holder saying it can be
```

distributed under the terms of this License. Such a notice grants a
world-wide, royalty-free license, unlimited in duration, to use that
work under the conditions stated herein. The "Document", below,
refers to any such manual or work. Any member of the public is a
licensee, and is addressed as "you". You accept the license if you
copy, modify or distribute the work in a way requiring permission
under copyright law.

A "Modified Version" of the Document means any work containing the
Document or a portion of it, either copied verbatim, or with
modifications and/or translated into another language.

A "Secondary Section" is a named appendix or a front-matter section of
the Document that deals exclusively with the relationship of the
publishers or authors of the Document to the Document's overall
subject (or to related matters) and contains nothing that could fall
directly within that overall subject. (Thus, if the Document is in
part a textbook of mathematics, a Secondary Section may not explain
any mathematics.) The relationship could be a matter of historical
connection with the subject or with related matters, or of legal,
commercial, philosophical, ethical or political position regarding
them.

The "Invariant Sections" are certain Secondary Sections whose titles
are designated, as being those of Invariant Sections, in the notice
that says that the Document is released under this License. If a
section does not fit the above definition of Secondary then it is not
allowed to be designated as Invariant. The Document may contain zero
Invariant Sections. If the Document does not identify any Invariant
Sections then there are none.

The "Cover Texts" are certain short passages of text that are listed,
as Front-Cover Texts or Back-Cover Texts, in the notice that says that
the Document is released under this License. A Front-Cover Text may
be at most 5 words, and a Back-Cover Text may be at most 25 words.

A "Transparent" copy of the Document means a machine-readable copy,
represented in a format whose specification is available to the
general public, that is suitable for revising the document
straightforwardly with generic text editors or (for images composed of
pixels) generic paint programs or (for drawings) some widely available
drawing editor, and that is suitable for input to text formatters or
for automatic translation to a variety of formats suitable for input
to text formatters. A copy made in an otherwise Transparent file
format whose markup, or absence of markup, has been arranged to thwart
or discourage subsequent modification by readers is not Transparent.
An image format is not Transparent if used for any substantial amount
of text. A copy that is not "Transparent" is called "Opaque".

Examples of suitable formats for Transparent copies include plain
ASCII without markup, Texinfo input format, LaTeX input format, SGML
or XML using a publicly available DTD, and standard-conforming simple
HTML, PostScript or PDF designed for human modification. Examples of
transparent image formats include PNG, XCF and JPG. Opaque formats
include proprietary formats that can be read and edited only by
proprietary word processors, SGML or XML for which the DTD and/or
processing tools are not generally available, and the
machine-generated HTML, PostScript or PDF produced by some word
processors for output purposes only.

The "Title Page" means, for a printed book, the title page itself,
plus such following pages as are needed to hold, legibly, the material

this License requires to appear in the title page. For works in formats which do not have any title page as such, "Title Page" means the text near the most prominent appearance of the work's title, preceding the beginning of the body of the text.

The "publisher" means any person or entity that distributes copies of the Document to the public.

A section "Entitled XYZ" means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as "Acknowledgements", "Dedications", "Endorsements", or "History".) To "Preserve the Title" of such a section when you modify the Document means that it remains a section "Entitled XYZ" according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy

```
a computer-network location from which the general network-using
public has access to download using public-standard network protocols
a complete Transparent copy of the Document, free of added material.
If you use the latter option, you must take reasonably prudent steps,
when you begin distribution of Opaque copies in quantity, to ensure
that this Transparent copy will remain thus accessible at the stated
location until at least one year after the last time you distribute an
Opaque copy (directly or through your agents or retailers) of that
edition to the public.

It is requested, but not required, that you contact the authors of the
Document well before redistributing any large number of copies, to
give them a chance to provide you with an updated version of the
Document.


4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under
the conditions of sections 2 and 3 above, provided that you release
the Modified Version under precisely this License, with the Modified
Version filling the role of the Document, thus licensing distribution
and modification of the Modified Version to whoever possesses a copy
of it.  In addition, you must do these things in the Modified Version:

A. Use in the Title Page (and on the covers, if any) a title distinct
   from that of the Document, and from those of previous versions
   (which should, if there were any, be listed in the History section
   of the Document).  You may use the same title as a previous version
   if the original publisher of that version gives permission.
B. List on the Title Page, as authors, one or more persons or entities
   responsible for authorship of the modifications in the Modified
   Version, together with at least five of the principal authors of the
   Document (all of its principal authors, if it has fewer than five),
   unless they release you from this requirement.
C. State on the Title page the name of the publisher of the
   Modified Version, as the publisher.
D. Preserve all the copyright notices of the Document.
E. Add an appropriate copyright notice for your modifications
   adjacent to the other copyright notices.
F. Include, immediately after the copyright notices, a license notice
   giving the public permission to use the Modified Version under the
   terms of this License, in the form shown in the Addendum below.
G. Preserve in that license notice the full lists of Invariant Sections
   and required Cover Texts given in the Document's license notice.
H. Include an unaltered copy of this License.
I. Preserve the section Entitled "History", Preserve its Title, and add
   to it an item stating at least the title, year, new authors, and
   publisher of the Modified Version as given on the Title Page.  If
   there is no section Entitled "History" in the Document, create one
   stating the title, year, authors, and publisher of the Document as
   given on its Title Page, then add an item describing the Modified
   Version as stated in the previous sentence.
J. Preserve the network location, if any, given in the Document for
   public access to a Transparent copy of the Document, and likewise
   the network locations given in the Document for previous versions
   it was based on.  These may be placed in the "History" section.
   You may omit a network location for a work that was published at
   least four years before the Document itself, or if the original
   publisher of the version it refers to gives permission.
K. For any section Entitled "Acknowledgements" or "Dedications",
   Preserve the Title of the section, and preserve in the section all
```

```
   the substance and tone of each of the contributor acknowledgements
   and/or dedications given therein.
L. Preserve all the Invariant Sections of the Document,
   unaltered in their text and in their titles.  Section numbers
   or the equivalent are not considered part of the section titles.
M. Delete any section Entitled "Endorsements".  Such a section
   may not be included in the Modified Version.
N. Do not retitle any existing section to be Entitled "Endorsements"
   or to conflict in title with any Invariant Section.
O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or
appendices that qualify as Secondary Sections and contain no material
copied from the Document, you may at your option designate some or all
of these sections as invariant.  To do this, add their titles to the
list of Invariant Sections in the Modified Version's license notice.
These titles must be distinct from any other section titles.

You may add a section Entitled "Endorsements", provided it contains
nothing but endorsements of your Modified Version by various
parties--for example, statements of peer review or that the text has
been approved by an organization as the authoritative definition of a
standard.

You may add a passage of up to five words as a Front-Cover Text, and a
passage of up to 25 words as a Back-Cover Text, to the end of the list
of Cover Texts in the Modified Version.  Only one passage of
Front-Cover Text and one of Back-Cover Text may be added by (or
through arrangements made by) any one entity.  If the Document already
includes a cover text for the same cover, previously added by you or
by arrangement made by the same entity you are acting on behalf of,
you may not add another; but you may replace the old one, on explicit
permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License
give permission to use their names for publicity for or to assert or
imply endorsement of any Modified Version.


5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this
License, under the terms defined in section 4 above for modified
versions, provided that you include in the combination all of the
Invariant Sections of all of the original documents, unmodified, and
list them all as Invariant Sections of your combined work in its
license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and
multiple identical Invariant Sections may be replaced with a single
copy.  If there are multiple Invariant Sections with the same name but
different contents, make the title of each such section unique by
adding at the end of it, in parentheses, the name of the original
author or publisher of that section if known, or else a unique number.
Make the same adjustment to the section titles in the list of
Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled "History"
in the various original documents, forming one section Entitled
"History"; likewise combine any sections Entitled "Acknowledgements",
and any sections Entitled "Dedications".  You must delete all sections
Entitled "Endorsements".
```

6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an "aggregate" if the copyright resulting from the compilation is not used to limit the legal rights of the compilation's users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document's Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled "Acknowledgements", "Dedications", or "History", the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document

except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, or distribute it is void, and will automatically terminate your rights under this License.

However, if you cease all violation of this License, then your license from a particular copyright holder is reinstated (a) provisionally, unless and until the copyright holder explicitly and finally terminates your license, and (b) permanently, if the copyright holder fails to notify you of the violation by some reasonable means prior to 60 days after the cessation.

Moreover, your license from a particular copyright holder is reinstated permanently if the copyright holder notifies you of the violation by some reasonable means, this is the first time you have received notice of violation of this License (for any work) from that copyright holder, and you cure the violation prior to 30 days after your receipt of the notice.

Termination of your rights under this section does not terminate the licenses of parties who have received copies or rights from you under this License. If your rights have been terminated and not permanently reinstated, receipt of a copy of some or all of the same material does not give you any rights to use it.

10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See http://www.gnu.org/copyleft/.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or any later version" applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation. If the Document specifies that a proxy can decide which future versions of this License can be used, that proxy's public statement of acceptance of a version permanently authorizes you to choose that version for the Document.

11. RELICENSING

"Massive Multiauthor Collaboration Site" (or "MMC Site") means any World Wide Web server that publishes copyrightable works and also provides prominent facilities for anybody to edit those works. A public wiki that anybody can edit is an example of such a server. A "Massive Multiauthor Collaboration" (or "MMC") contained in the site means any set of copyrightable works thus published on the MMC site.

"CC-BY-SA" means the Creative Commons Attribution-Share Alike 3.0 license published by Creative Commons Corporation, a not-for-profit corporation with a principal place of business in San Francisco, California, as well as future copyleft versions of that license published by that same organization.

"Incorporate" means to publish or republish a Document, in whole or in

```
part, as part of another Document.

An MMC is "eligible for relicensing" if it is licensed under this
License, and if all works that were first published under this License
somewhere other than this MMC, and subsequently incorporated in whole or
in part into the MMC, (1) had no cover texts or invariant sections, and
(2) were thus incorporated prior to November 1, 2008.

The operator of an MMC Site may republish an MMC contained in the site
under CC-BY-SA on the same site at any time before August 1, 2009,
provided the MMC is eligible for relicensing.


ADDENDUM: How to use this License for your documents

To use this License in a document you have written, include a copy of
the License in the document and put the following copyright and
license notices just after the title page:

    Copyright (c)  YEAR  YOUR NAME.
    Permission is granted to copy, distribute and/or modify this document
    under the terms of the GNU Free Documentation License, Version 1.3
    or any later version published by the Free Software Foundation;
    with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts.
    A copy of the license is included in the section entitled "GNU
    Free Documentation License".

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts,
replace the "with...Texts." line with this:

    with the Invariant Sections being LIST THEIR TITLES, with the
    Front-Cover Texts being LIST, and with the Back-Cover Texts being LIST.

If you have Invariant Sections without Cover Texts, or some other
combination of the three, merge those two alternatives to suit the
situation.

If your document contains nontrivial examples of program code, we
recommend releasing these examples in parallel under your choice of
free software license, such as the GNU General Public License,
to permit their use in free software.
```

E.2 LGPL: GNU Lesser Gnu Public License

```
                   GNU LESSER GENERAL PUBLIC LICENSE
                       Version 3, 29 June 2007

 Copyright (C) 2007 Free Software Foundation, Inc. <http://fsf.org/>
 Everyone is permitted to copy and distribute verbatim copies
 of this license document, but changing it is not allowed.


  This version of the GNU Lesser General Public License incorporates
the terms and conditions of version 3 of the GNU General Public
License, supplemented by the additional permissions listed below.

  0. Additional Definitions.

  As used herein, "this License" refers to version 3 of the GNU Lesser
```

```
General Public License, and the "GNU GPL" refers to version 3 of the GNU
General Public License.

  "The Library" refers to a covered work governed by this License,
other than an Application or a Combined Work as defined below.

  An "Application" is any work that makes use of an interface provided
by the Library, but which is not otherwise based on the Library.
Defining a subclass of a class defined by the Library is deemed a mode
of using an interface provided by the Library.

  A "Combined Work" is a work produced by combining or linking an
Application with the Library.  The particular version of the Library
with which the Combined Work was made is also called the "Linked
Version".

  The "Minimal Corresponding Source" for a Combined Work means the
Corresponding Source for the Combined Work, excluding any source code
for portions of the Combined Work that, considered in isolation, are
based on the Application, and not on the Linked Version.

  The "Corresponding Application Code" for a Combined Work means the
object code and/or source code for the Application, including any data
and utility programs needed for reproducing the Combined Work from the
Application, but excluding the System Libraries of the Combined Work.

  1. Exception to Section 3 of the GNU GPL.

  You may convey a covered work under sections 3 and 4 of this License
without being bound by section 3 of the GNU GPL.

  2. Conveying Modified Versions.

  If you modify a copy of the Library, and, in your modifications, a
facility refers to a function or data to be supplied by an Application
that uses the facility (other than as an argument passed when the
facility is invoked), then you may convey a copy of the modified
version:

   a) under this License, provided that you make a good faith effort to
   ensure that, in the event an Application does not supply the
   function or data, the facility still operates, and performs
   whatever part of its purpose remains meaningful, or

   b) under the GNU GPL, with none of the additional permissions of
   this License applicable to that copy.

  3. Object Code Incorporating Material from Library Header Files.

  The object code form of an Application may incorporate material from
a header file that is part of the Library.  You may convey such object
code under terms of your choice, provided that, if the incorporated
material is not limited to numerical parameters, data structure
layouts and accessors, or small macros, inline functions and templates
(ten or fewer lines in length), you do both of the following:

   a) Give prominent notice with each copy of the object code that the
   Library is used in it and that the Library and its use are
   covered by this License.

   b) Accompany the object code with a copy of the GNU GPL and this license
   document.
```

```
  4. Combined Works.

  You may convey a Combined Work under terms of your choice that,
taken together, effectively do not restrict modification of the
portions of the Library contained in the Combined Work and reverse
engineering for debugging such modifications, if you also do each of
the following:

   a) Give prominent notice with each copy of the Combined Work that
   the Library is used in it and that the Library and its use are
   covered by this License.

   b) Accompany the Combined Work with a copy of the GNU GPL and this license
   document.

   c) For a Combined Work that displays copyright notices during
   execution, include the copyright notice for the Library among
   these notices, as well as a reference directing the user to the
   copies of the GNU GPL and this license document.

   d) Do one of the following:

       0) Convey the Minimal Corresponding Source under the terms of this
       License, and the Corresponding Application Code in a form
       suitable for, and under terms that permit, the user to
       recombine or relink the Application with a modified version of
       the Linked Version to produce a modified Combined Work, in the
       manner specified by section 6 of the GNU GPL for conveying
       Corresponding Source.

       1) Use a suitable shared library mechanism for linking with the
       Library.  A suitable mechanism is one that (a) uses at run time
       a copy of the Library already present on the user's computer
       system, and (b) will operate properly with a modified version
       of the Library that is interface-compatible with the Linked
       Version.

   e) Provide Installation Information, but only if you would otherwise
   be required to provide such information under section 6 of the
   GNU GPL, and only to the extent that such information is
   necessary to install and execute a modified version of the
   Combined Work produced by recombining or relinking the
   Application with a modified version of the Linked Version. (If
   you use option 4d0, the Installation Information must accompany
   the Minimal Corresponding Source and Corresponding Application
   Code. If you use option 4d1, you must provide the Installation
   Information in the manner specified by section 6 of the GNU GPL
   for conveying Corresponding Source.)

  5. Combined Libraries.

  You may place library facilities that are a work based on the
Library side by side in a single library together with other library
facilities that are not Applications and are not covered by this
License, and convey such a combined library under terms of your
choice, if you do both of the following:

   a) Accompany the combined library with a copy of the same work based
   on the Library, uncombined with any other library facilities,
   conveyed under the terms of this License.
```

```
   b) Give prominent notice with the combined library that part of it
   is a work based on the Library, and explaining where to find the
   accompanying uncombined form of the same work.

  6. Revised Versions of the GNU Lesser General Public License.

  The Free Software Foundation may publish revised and/or new versions
of the GNU Lesser General Public License from time to time. Such new
versions will be similar in spirit to the present version, but may
differ in detail to address new problems or concerns.

  Each version is given a distinguishing version number. If the
Library as you received it specifies that a certain numbered version
of the GNU Lesser General Public License "or any later version"
applies to it, you have the option of following the terms and
conditions either of that published version or of any later version
published by the Free Software Foundation. If the Library as you
received it does not specify a version number of the GNU Lesser
General Public License, you may choose any version of the GNU Lesser
General Public License ever published by the Free Software Foundation.

  If the Library as you received it specifies that a proxy can decide
whether future versions of the GNU Lesser General Public License shall
apply, that proxy's public statement of acceptance of any version is
permanent authorization for you to choose that version for the
Library.
```

Chapter 6

Index

A
API, 7, 32
attributes, 33, 35, 36
    data, 2, 10, 13
    global, 13
automatic plotting, *see* NeXus basic motivation, default plot
axes, 13, 30
axis, 13, 30

C
classes
    base classes
        NXdata, 3
        NXentry, 3
        NXinstrument, 3
        NXsample, 3
coordinate systems, 16
    CIF, 16
    IUCr, 17
    McStas, 16, 17
    NeXus polar coordinate, 16
    spherical polar, 17
    transformations, 18
        order of operations, 18

D
data objects
    attributes, 2, 13
        global, 13
    data items, *see* fields
    fields, 2, 12, 13
    groups, 2, 12
date and time, 13, 29
dimension, 13, 30, 36
    data set, 31, 44, 67
    dimension scales, 30, 31
    fastest varying, 30
    storage order, 26
dimension scale, 13, 15, 16, 32

E
enumeration, 30
example
    simple, 6
    very simple, 4, 5

F
FAQ, 66
file
    browse, 10
    read, 10
    write, 9

G
geometry, 16, 17

H
HDF, 32, 56, 60
    Scientific Data Sets, *see* data objects, fields
hierarchy, 2, 12, 19, 21, 40, 41, 63
    example NeXus file, 2
    NeXus, 47

I
installation, 56
instrument definitions, 5
ISO 8601, *see* date and time
issue reporting, *see* TRAC

L
License
    FDL, 68
    LGPL, 75
link, 2, 14, 30, 36, 62, 67

M
mailing lists, 51
McStas, 16
metadata, 40, 41, 47, 59, 63
monitor, 32

N
NAPI, 1, 8, 9, 48, 63, 67
    bypassing, 32
NAPIlink, 36
NeXus, 1
    Community, 50

Design Principles, 2
low-level file formats, 32
NeXus basic motivation, 6
default plot, 2, 3, 7, 9, 13, 15, 21, 31, 64, 67
defined dictionary, 8
unified format, 1, 7
NeXus International Advisory Committee, *see* NIAC
NIAC, 50, 51, 66
NXdata, 30, 31
NXDL, 14, 62, 67

R
rank, 9, 13, 31, 32, 44
roadmap, 54
rules, 1, 61
HDF, 12
HDF5, 25
naming, 14, 19
NeXus, 19, 59–62
naming, 25
NXDL, 60, 61
Schematron, 62

S
Schematron, 61, 62
Scientific Data Sets, *see* data objects, fields
SDS, *see* data objects, fields
subversion, 52

T
target, *see* link
time, *see* date and time
TRAC, 54
tutorial
WONI, 37

U
UDunits, 30
units, 2, 10, 13, 30
utility
dave, 64
gda, 64
gumtree, 64
HDF5 Command Line Tools, 64
hdfexplorer, 65
hdfview, 65
idl, 65
IgorPro, 65
isaw, 64
lamp, 64
mantid, 64
matlab, 65
nexpy, 64
nxbrowse, 10, 63
nxconvert, 63
nxdir, 63
nxdump, 64
nxingest, 63
nxplot, 64
nxsummary, 63
nxtranslate, 63
NXvalidate, 64
nxvalidate, 63
opengenie, 64
pymca, 64

V
validation, 59, 61
NeXus data files, 61
NXDL rules, 62
NXDL specifications, 61
XSLT files, 62
verification, *see* validation

W
WONI, 37

X
XML, 7, 56
XML Schema (XSD), 61, 62
XSLT, 61, 62